The Verification Experts

SYNC SERIES

GPS Disciplining & Holdover for Field Testing

Introduction, Examples, Analysis and Recommendations

By Ildefonso M. Polo
Dir. Product Marketing – Transport & Synchronization

December 2016 | Rev. B00

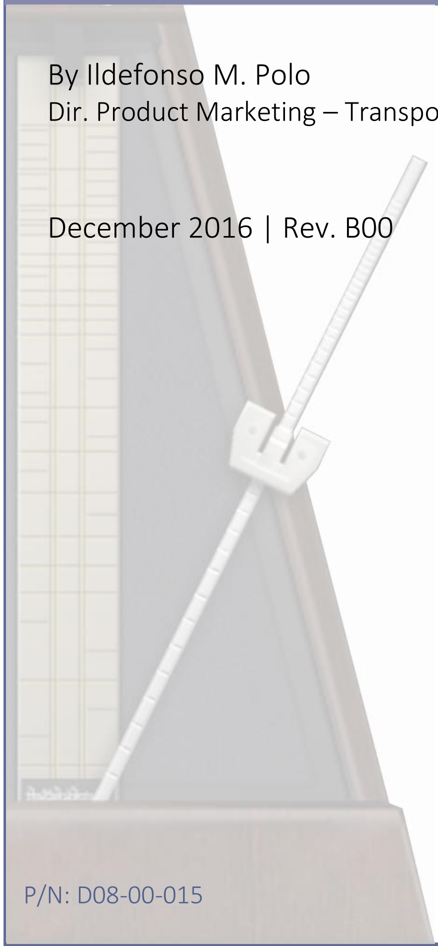

P/N: D08-00-015

Notice:

The data for this article was acquired using VeEX's TX300S and/or RXT-1200 Test Platforms with GNSS and Atomic Clock options, and Multi-Service Test Modules. TE data is unfiltered and graphs were captured with VeEX Wander Analysis PC Software and/or built-in Clock Wander Analysis function.

For more technical resources, visit the VeEX Inc. website at www.veexinc.com.

For assistance or questions related to the use of VeEX products or its features, functions and capabilities, please contact our local representative or contact VeEX directly by phone or e-mail.

VeEX Inc.
Phone: +1 510 651 0500
E-mail: info@veexinc.com
Website: www.veexinc.com

Notes to those who may not be familiar with the TIE or TE wander graphs used in this article:

- Each data point in the blue line represents the instantaneous phase error (or time error) over time.
- The slope (tangent) at any point in the graph represents the frequency offset at that point in time.
- Any exponential trend represents frequency drift.
- The vertical (time error) scale is in nanoseconds while the horizontal (test time) is in seconds.

Using GPS Disciplining & Holdover for Field Testing

Introduction, Examples, Analysis and Recommendations

One of the challenges synchronization engineers and technicians often face in the field, is the lack of proper Frequency and Timing (Phase) references required to validate 1PPS clocks' Accuracy and Stability at remote sites. This leaves them with one practical choice: The use of GNSS-disciplined oscillators (GPSDO) as reference clocks. "Extended" Phase Holdover has also been proposed as an alternative reference to help with indoors testing (no GNSS signal). So, it is important to understand how disciplining and the "extended holdover" work, their limitations and how to use them properly, to have a clear and practical idea of what to expect and what not.

GNSS, GPS in particular, is the best (if not the only practical) option available to get accurate time, timing and frequency almost anywhere in the world. But it is not perfect. Knowing the precautions and GNSS limitations help in defining correct test procedures and setting realistic expectation for field testing.

This article includes simple explanations on how disciplining and holdover work, so it benefits a greater audience, including field engineers and technicians who are new to synchronization. Nevertheless, experienced individuals should not be discouraged from reading through, since some of the information bits may "strike a note" and/or encourage participation in discussions and contribute to improve the little information currently available.

Basic Disciplining Introduction

Basically, a GNSS-disciplined Clock consists of a high-quality (precision) oscillator that is continuously being corrected using the coordinated universal timing signal (UTC), or standard second, in the form of 1PPS pulses recovered from the GPS signals. In theory, the rising edge of the One-Pulse-Per-Second (1PPS) represents the beginning of a standard second (in practice it is very close).

Frequency and Phase are tightly related, but don't assume that having accurate frequency guarantees accurate phase. In that sense, they are quite different. A disciplining control system will actually make its frequency source slightly inaccurate, on purpose, in order to achieve and maintain accurate phase. There is also noise, jitter, wander and temperature dependence in an otherwise perfect frequency source, which cause random (phase) walk.

Relationship between Frequency and Phase

A frequency of 1GHz (1ns cycle) is used in this simplified example to make mental and visual calculations easier. The illustration below depicts the relationship between Frequency Offset and Phase Drift, which is used in the 1PPS disciplining process to control the position of the output pulse, relative to the standard second. The darker cycles indicate the end of the billionth cycles.

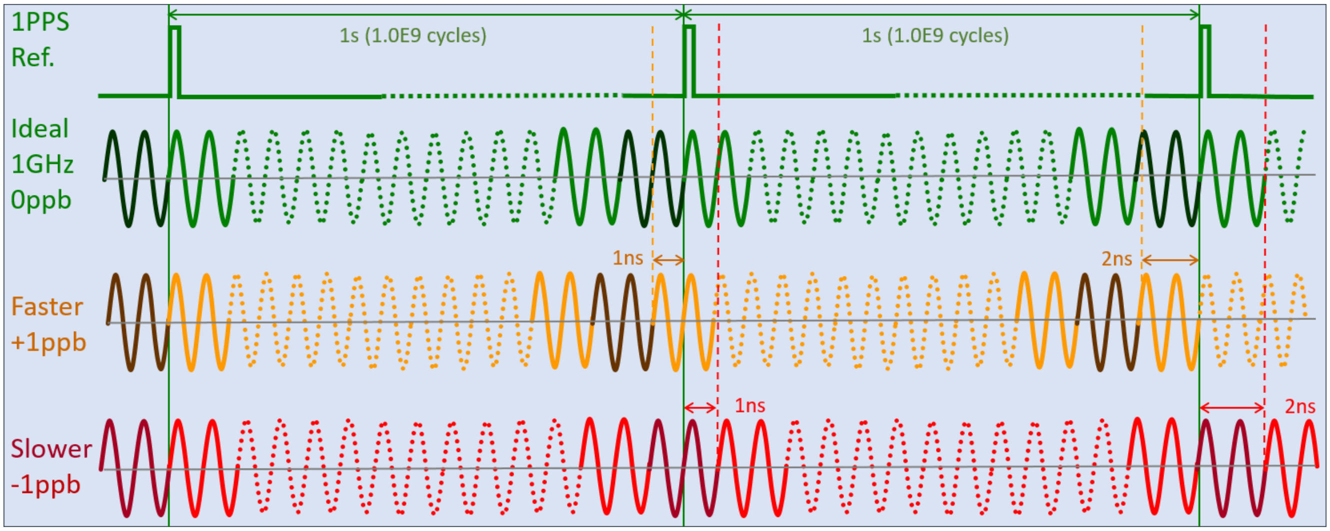

If a 1,000,000,000 counter is used to count one second and the frequency is 1,000,000,001Hz (+1 ppb), there is one extra cycle (≈1ns), which would make the counter roll over earlier, adding about 1ns phase error for every second that goes by (cumulative). It is similar for 999,999,999Hz (-1 ppb) as one cycle (≈1ns) will be missing from every second, making the counter roll over about 1ns later, then 2ns, 3ns, ... That is, positive frequency offset creates positive time drift and negative frequency offset creates negative time drift, at a rate of X ppb = X ns/s.

Frequency Disciplining (Syntonization)

To correct any frequency deviation (offset) in the local oscillator, the system could basically count the number of cycles that occur between the rising edges of two consecutive 1PPS pulses (i.e. "exactly" one second). If it is a 10MHz oscillator and 10,000,010 cycles are counted, then the oscillator is running 1 ppm (1.0E-6) too fast. Then the control system would change (steer) the oscillator's frequency to slow it down. It counts again and gets 10,000,001 cycles, indicating that the oscillator is now running with 0.1 ppm offset and uses that information to make another correction to perhaps get 9,999,999.9 (-0.01 ppm). Then finer and finer corrections are made, until the count is as close to 10,000,000.00000 cycles as the system can measure, hence achieving a very high frequency accuracy (e.g. in the order of parts-per-trillion, or 1.0E-12). It continues making those very small corrections to compensate for any changes due to ambient temperature variation and oscillator's ageing. Basically, the oscillator is being actively "soft" calibrated in real time and it should now be highly accurate.

One of the caveats here is that the timing signals recovered from GNSS are not ideal. In real life those faint radio signals have to deal with our imperfect world (e.g. ionospheric and atmospheric changes, multi-path, interference, electronics, temperature, etc.). That otherwise "perfect" second varies (wanders) over time. It means that some of the frequency measurements made by the disciplining control system would be a bit off and some slightly wrong correction calculations would be made. Disciplining systems mitigate this problem by applying some statistical algorithms before issuing the steer correction to the oscillator. Having a high-quality oscillator technology would also help as they are difficult to steer in the short term, creating a dampening effect that helps filter that noise (errors). The end result is a frequency source that benefits from the best of both worlds: The short-term stability of a good oscillator plus the long-term accuracy of GNSS, resulting in a highly accurate and stable frequency source.

Phase/Timing Disciplining (Synchronization)

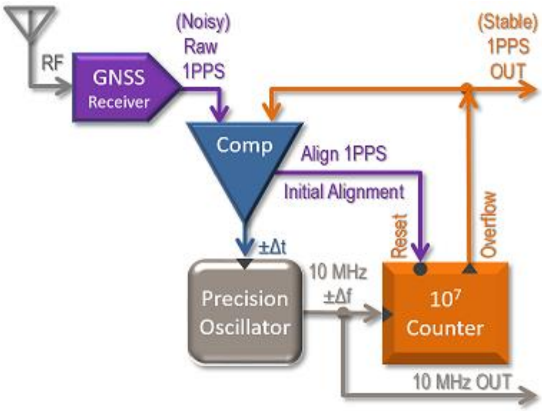

The process is very similar to the older frequency-only GPS Clocks described earlier, but in this case the priority is to keep the 1PPS phase aligned to the standard second. That is, the control system can intentionally change the frequency of the oscillator to keep the 1PPS phase aligned to the standard second (UTC). As one could imagine, some frequency purists may not like this approach. Imagine starting with a perfectly calibrated 10MHz oscillator and the first thing the disciplining control system does is to deliberately change its frequency to make the 1PPS pulse drift closer towards its ideal position.

If the oscillator's 10MHz disciplined output described earlier is fed to a 10 million counter, it would produce an overflow pulse every time it rolls over from 9,999,999 to zero. That is, every second. Initially the position of the resulting 1PPS output would be arbitrary. The first task for the phase disciplining control system is to perform a rough alignment, using the GPS 1PPS pulse to reset the counter. This leaves the 1PPS output within 100ns (one 10MHz cycle) from the standard second. To fine tune the 1PPS phase alignment to UTC, the control system has to slowly change the oscillator's frequency. For example, if the local 1PPS is 30ns late, the frequency of the oscillator is slightly increased to count faster and start reducing the phase error (absolute time error). On the other hand, if the local 1PPS happens before UTC (e.g. 25ns earlier), the oscillator is slowed down to delay the counter. It could be said that the amount of new frequency offset corrections are proportional to the remaining time error (plus some amount of vendor-specific statistical magic). It means that the local phase should slowly converge towards UTC and the amount of frequency offset (inaccuracy) in the output also reduces. Once this convergence happens, the GNSS Clock should be outputting fairly accurate and stable 1PPS timing and 10MHz frequency references.

In the absence of any other references or instruments to confirm accuracy and stability, field users should monitor the ±ΔT value (input vs. output) to identify when the disciplining has converged (horizontal trend at 0ns). At that point, the test set should have reached the best accuracy it could possibly get under the existing conditions.

Why don't just use the 1PPS directly from the GNSS/GPS Receiver?

- First, it is quite noisy and needs to be stabilized.
- Second, it could go away at any time, so a good local oscillator is required to help bridge any outages (holdover).

Disciplining for Synchronization Testing in the Field

An instrument's internal precision clock can be disciplined by:

a) Connecting it to a highly accurate and stable 1PPS, from a traceable Cs or Rb timing reference (e.g. PRTC).

b) Using the 1PPS from a built-in GNSS receiver.

Option (a) is much faster, more accurate and should provide longer and more predictable holdover. But, if we had access to PRTCs in the field, we would use them for testing, directly. There would be no need for extended holdover. Nonetheless, direct disciplining can come handy in central offices, labs, data centers, etc., where there is a need to quickly transport accurate timing within the premises, to avoid long cables (cables add phase error). So, we are left with option (b) and the rest of this article focuses on GNSS disciplining for field testing.

Using GNSS Disciplined Clocks for Testing

High quality GNSS Clocks are excellent timing and frequency references for:

- Frequency accuracy (Offset) and stability measurements (Wander).
- 1PPS phase accuracy (absolute Time Error) and stability measurements (Wander).
- One-Way Delay (Latency) and Asymmetry measurements.

Keep in mind that not all GPS Receivers are necessarily good GPS Clocks. There are different qualities of GPS Clocks using Cs, Rb, DOCXO, OCXO or even TCXO oscillators. The oscillator's quality plays a key role in the GNSS Clock performance required for test and measurements. Atomic Clock technology is preferred for test and measurement.

When performing frequency-related tests, such as wander (TIE, MTIE, TDEV), cable lengths may have not been an issue. With Time Error the situation is completely different. Every meter of cable introduces about 5ns delay (error) which can add up very quickly.

- The antenna cable length must be properly documented and programmed into any stationary GNSS Clock for 1PPS delay compensation.
- Test cables (from DUT and 1PPS reference) must be of the same length to cancel out their respective delays.

Holdover Introduction

Although GNSS signals "rain" from the sky, they are quite faint, have poor penetration and are easily obscured. They are not available (or reliable) indoors, in "urban canyons" or noisy RF environments. So being able to hold the last known frequency and timing (phase) during a GNSS outage becomes a necessary requirement. This is why GNSS receivers are paired with high-quality oscillators, to create clock sources with holdover characteristics that can bridge temporary RF outages.

Unfortunately, those obscured places are the ones we need to run tests at (shielded cabinets, containers, base stations, central offices, equipment rooms, office buildings, basements, etc.). The proposed solution is to extend the use of holdover into becoming a temporary mean of transporting time synchronization from outdoors to indoors.

Before jumping into the holdover-for-testing explanation, keep in mind that the phase disciplining process constantly changes the oscillator's frequency to keep its 1PPS aligned to the standard UTC second. Even though those fine adjustments are in the order of parts-per-trillion, they will affect the final holdover performance.

Using Disciplining and Holdover in Actual Field Test Applications

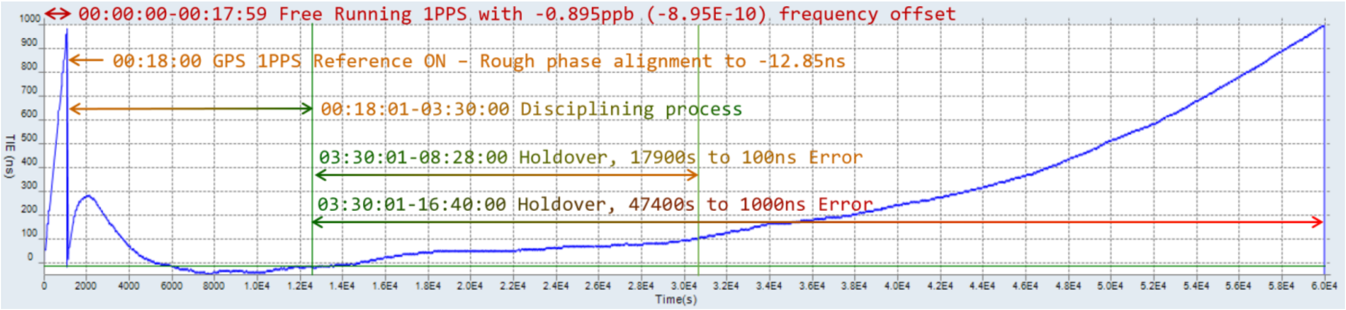

This graph depicts the full Time Error (TE) cycle of a test set's internal 1PPS reference, starting from free running, disciplining, to holdover. A relatively short holdover behavior was selected as an example, so all the different phases can be easily visualized.

The test was ran in conditions similar to what field technicians would encounter at remote sites. The main difference is that we had access to a 1PPS reference (PRTC) and equipment to monitor and document the internal 1PPS behavior for this article, but field technicians won't have the same luxury. That is the reason why it is important to carry reliable test instruments that provide as much information about their internal status and readiness as possible.

A test set, equipped with GPS receiver and chip-scale Atomic Clock oscillator, was selected for the experiment. It was turned on and its internal GPS-disciplined 1PPS reference was constantly recorded, but that extra TE information was not used for any decision making during the test. Let's interpret this graph one section at a time.

Free Running Stage

The initial part of the trace shows the internal 1PPS phase compared to the reference, while the GPS disciplining function was still disabled (free running).

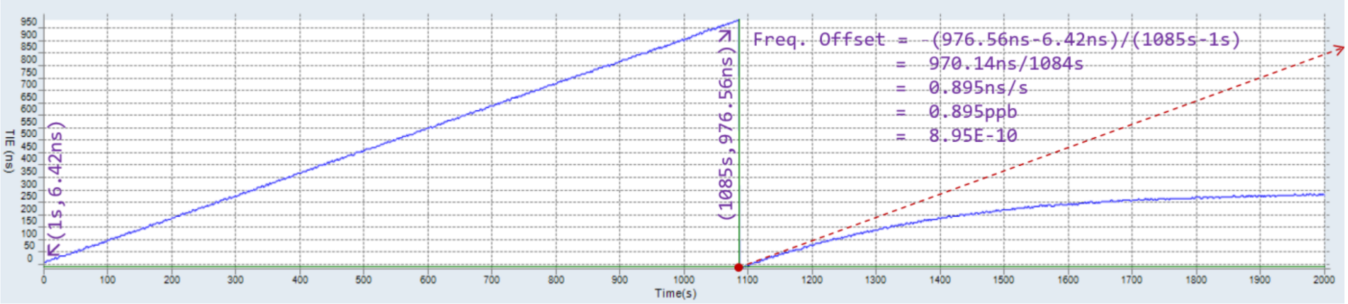

The first 1,085 seconds of the TE trace indicate that the oscillator is very stable (straight line) but it has a free-run accuracy or frequency offset (ramp trend) of 0.895 ppb, which causes the phase to drift at a rate of 0.895 ns/s. The straight TE lines, which identify highly stable oscillators, make it easier to measure its frequency offset by using simple math to calculate the slope (ΔY/ΔX or ΔTE/Δt). Modern test sets can also performs the offset calculation on their own, with the advantage that they can do it for more complex scenarios.

Disciplining Process

18 minutes later (1,085s) the GPS (already locked to satellites) is enabled, to get the disciplining process started. The first thing that happens is the internal 1PPS being forced to align with the 10MHz cycle closest to the UTC second. In this case, the TE reads -12.5ns.

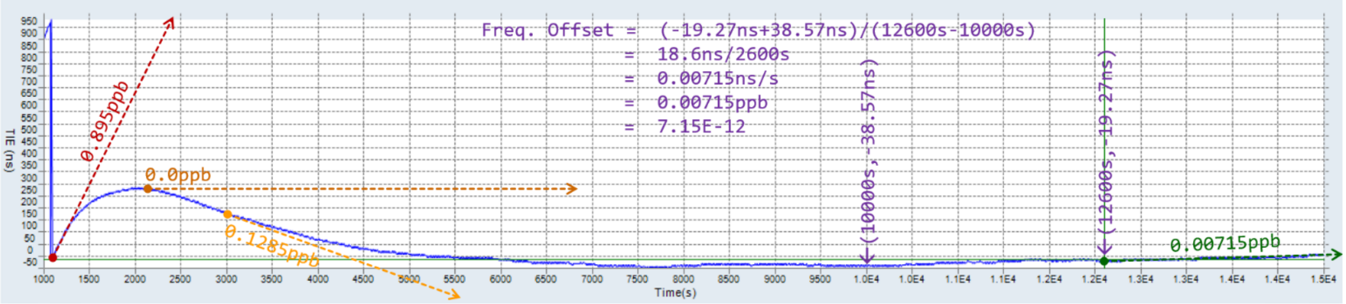

The first dot shows a trend indicating that the internal oscillator still wants to run at 10,000,000.00895Hz (10MHz +0.895ppb), after the initial rough phase correction. The disciplining control system starts to steer the oscillator into running slower and at the 2,000s mark it reaches exactly 10,000,000Hz (0.0ppb). If the goal was to achieve frequency accuracy only, the process could have been stopped at this point and let it stay in that perfectly horizontal trend indicated by the dotted line. But the priority is Phase (Time) alignment and 300ns error has accumulated in the process of getting to this point, which still need to be corrected. So, the process continues by making the oscillator run a little bit slower than its ideal frequency, in order to make the pulse's phase drift to the "left", to reduce the phase error (as described in the initial diagram).

The third dot (at 3,000s) shows the oscillator running at 9,999,999.998715Hz (-0.1285ppb). As the 1PPS phase error approaches 0.0ns, the oscillator starts to run a bit faster to make it converge. Depending on the conditions and settings, the whole process may take long time, which could be an issue for field testing, if not done correctly.

The fourth dot (at 12,600s or 03:30:00) shows the oscillator running at around 10,000,000.0000715Hz (0.00715ppb), before the GPS 1PPS reference was disconnected. That is two orders of magnitude better than its free-running frequency. At first, this may look impressive, but even in a controlled lab environment a 0.00715 ns/s phase drift rate would cause the accumulation of 100ns error in about 13,986s (03:53:06).

Also note that, after checking the internal phase graph to confirm that the disciplining process has stabilized, the point of disconnection is basically selected at random (by a human). No one can really guarantee what the exact residual frequency offset is at every single occasion. It could be slightly larger or smaller, positive or negative, so there is a degree of uncertainty in the resulting holdover behavior. Keep in mind that field technicians don't have access to some of the measurements or calculations shown here, when making the decision to force the test set (or external reference) into holdover. So, having clear information about the internal disciplining process and status becomes very important.

Forcing the Test Set into Holdover

At the 12,600s mark the GPS was disabled, to force the test set's internal 1PPS reference into holdover mode, so one can enter the premises to perform Absolute Time Error, Wander or One-Way Latency measurements.

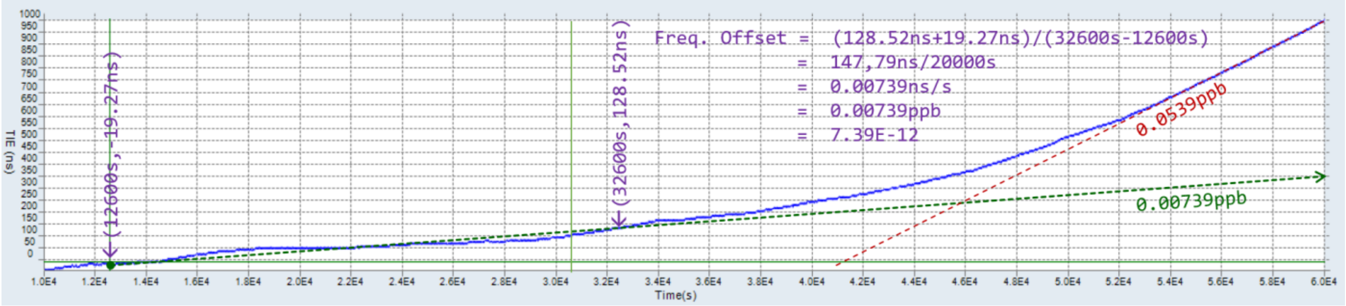

At 0.00739 ns/s phase drift rate, the initial trend is very close to what was calculated before the holdover started. This confirms that the oscillator will keep the last frequency correction to try holding the last known time. The first 100ns error are accumulated after 17,900s (04:58:20) of holdover and the 1,000ns threshold is reached after 47,400s (13:19:00). Note that the oscillator's frequency starts to drift a bit at the end of the day (5:45pm or 10:19am+34,000s), due to significant drop in ambient temperature (winter). Nonetheless, it covered the regular working hours. This effect would be less noticeable in equipment rooms with continuous air conditioning (HVAC).

Repeatability

GPS-disciplined Frequency references have an advantage, their frequency correction factor (steer) can be measured, stored and re-applied later on, as long as the oscillator's frequency retrace remains the same every time it is turned on. That is not the case for Phase references. They must remain powered in order to keep time. The selected phase reference must be battery powered in order to be used in field applications.

It is also important to know how the selected portable or built-in GNSS-disciplined Clock Source performs, under the conditions specific to a network and its environment. As mentioned before, the total holdover time to some amount of allowed uncertainty (or error) is mainly determined by (a) the frequency offset of the oscillator at the moment the GPS 1PPS is disabled or disconnected, (b) the stability of the oscillator, (c) temperature variations and (d) phase noise. Once a good disciplined oscillator has been properly steered into parts-per-trillion accuracy, the point at which the GPS is disconnected could be one of the most important factors in determining how long the holdover would be. It is all about the frequency offset and temperature stability from that point on. Each time the same procedure is ran, the results could be slightly different.

A week later, the same test set was used to run the experiment again, using approximately the same times to start the disciplining process and force the holdover. The graph shows the differences between the first (blue) and second (red) tests. Temperature and sky condition were slightly different since the second test was performed on a rainy day. That is good example of what real-life field conditions mean.

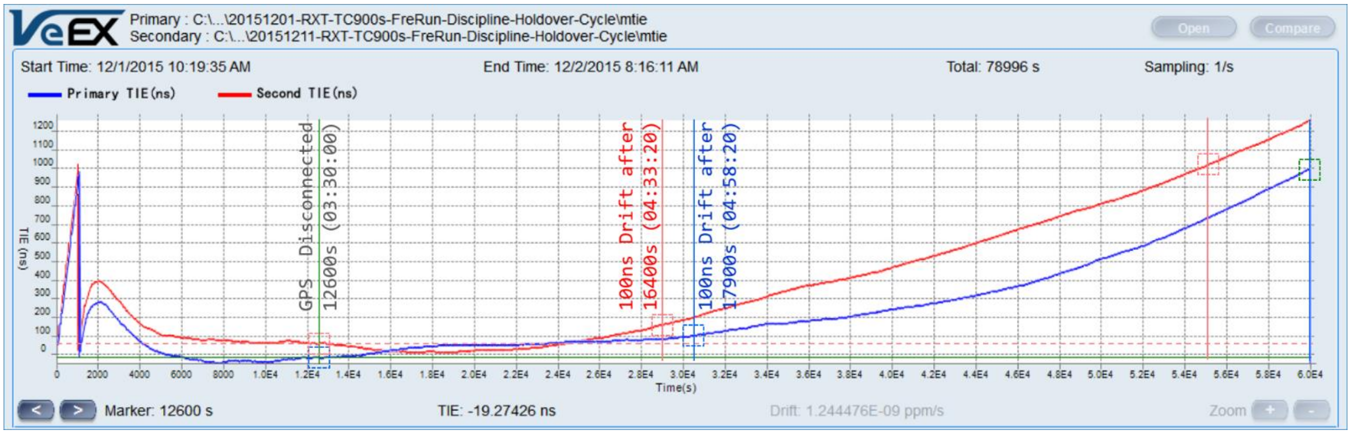
Primary : C:\...\20151201-RXT-TC900s-FreRun-Discipline-Holdover-Cycle\mtie
Secondary : C:\...\20151211-RXT-TC900s-FreRun-Discipline-Holdover-Cycle\mtie

Start Time: 12/1/2015 10:19:35 AM
End Time: 12/2/2015 8:16:11 AM
Total: 78996 s
Sampling: 1/s

Marker: 12600 s
TIE: -19.27426 ns
Drift: 1.244476E-09 ppm/s

As expected, the two traces are slightly different. There is a baseline offset of about 70ns between the original and the new TE traces, which could be due to the internal GPS Clock plus the external GPS-disciplined reference uncertainties. Also, in the new test results (red) the first 100ns accumulate faster (25 minutes earlier) and reaches the 1,000ns mark earlier, after 42,000s of holdover.

If 100ns uncertainty is considered acceptable, it could be said that such holdover reference could be used for up to four hours in both cases. It is important for field users to be aware that disciplined oscillators don't behave exactly the same every time. At the same time one can't assume that holdover performance is always that short. The better you know the basic theory and the instruments, the more reliable the holdover performance would be. If the same test is ran again, carefully applying all the guidance provided in this document, the performance can improve significantly, as shown by the results below.

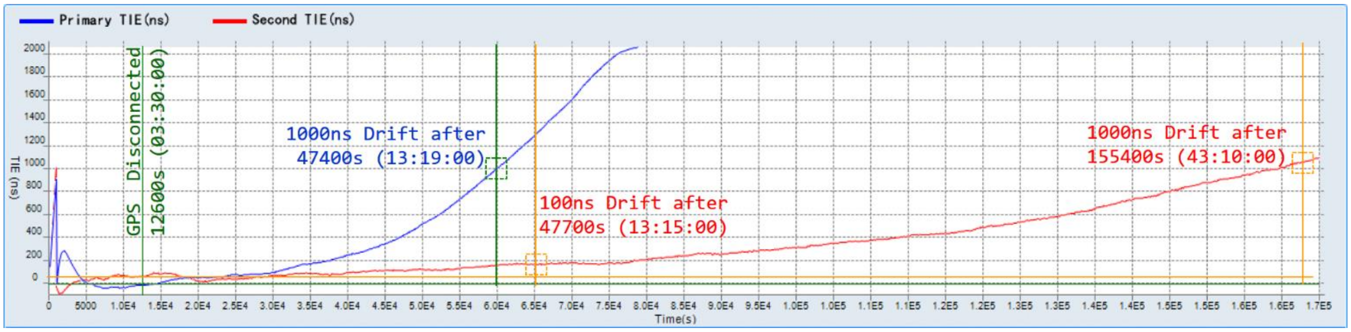

In this last example (red), the test set runs for more than 13 hours before accumulating 100ns error and it takes more than 43 hours to accumulate 1,000ns error. Significantly better than the original (blue) one, and it was also a field test (not a controlled lab test). As we all learn from true field experiences, T&M vendors may be able to automate part of the holdover initiation process to guarantee the best holdover possible under given conditions.

Can Phase Drift Effects, Due to Frequency Offset, be Corrected to Improve Repeatability?

It would be technically possible. The linear phase drift effects (stable frequency offset) can be measured with great precision and mathematically removed, but it could be cumbersome and/or too expensive for field testing.

- Some experts suggest having a GPS-disciplined time reference in the vehicle, which could be used at the beginning and at the end of each test to measure the instrument's internal 1PPS drift, calculate the frequency offset associated with it and mathematically removing its contribution from the recorded data. This could be an expensive proposition and it requires users to run outside after every test to measure the error accumulated by the portable reference.
- A simpler approach may be to make users go outside, after every test, and resync to GPS so the instrument can measure its accumulated phase error, calculate its phase drift rate and apply a (linear) mathematical correction to the recorded TE data. It's technically possible today, but the running outside part still needs some thoughts.

Portability and Battery Operation

The use or need for extended holdover in test and measurement implies battery backup, at least to keep the internal clock ticking while the instrument is being transported to the test site. Since holdover degrades timing accuracy over time, it is best if the test set can be disciplined on site (e.g. via GPS) to minimize transport time and optimize accuracy.

Another factor that makes portability a key factor are the test sites themselves. Time synchronization is being driven by wireless technologies such as LTE-TDD and LTE Advanced, and the test points are usually in places that may not be so easy to reach. eNB, containers, street cabinets, rooftops and towers, among others, would require full battery operation, small size, low weight and maneuverability. Basically the oscillator technology used defines the amount of power required for warm up and continuous operation, and the autonomy required to run the tests would define the amount of batteries, weight and size of the test gear. Traditional Cs and Rb can be very accurate and stable, but require a lot of power to warm up and operate, limiting their applicability to field testing. Newer technologies, such as chip-scale Atomic Clocks, aim to reduce the size of the physics package and required components, to minimize power consumption and enable full field operation.

Ideally, a synchronization verification test instrument should be able to be disciplined, transport timing, and run the required measurements on site, even if access to AC or DC power is not available or limited.

Notes about Using Forced Holdover for Testing

As a general rule, the use of phase holdover for testing should only be considered as an alternative, for those occasions in which GPS coverage is not available or reliable. If GPS signal or any other 1PPS reference are available on site, use them.

In those particular examples presented earlier, one could say that the test set would be able to perform fairly accurate Absolute Time Error and Wander measurement within four hours of entering holdover mode (assuming that 100ns uncertainty is acceptable), and One-Way-Delay test could be run for at least eight hours (assuming 1µs accuracy). But that may not always be the case. Results may vary depending on the disciplining settings.

Don't forget that the GPS receiver used to discipline the oscillator also has its own phase error, which may vary depending on atmospheric and ionospheric conditions, as well as antenna placement, interference and temperature variation. So there is an uncertainty to start with (<100ns per ITU-T G.8272 clauses 6.1 and 6.2).

Another issue to consider is the amount of time the GNSS Clock takes to discipline, in order to achieve accurate and stable frequency and phase alignment. Three hours may be considered too much of a wait for a field test, which may be impractical for some applications. There are ways to make it synchronize faster, but not without some tradeoffs. This is why it is important to fully understand the concept, so practical guidance, procedures and fair expectations are put in place to achieve the desired goals. Here are some suggestions:

- Adjust the GPS Clock's Time Constant to make the disciplining control system react faster and converge quicker. But, this may affect the final accuracy of the 10MHz and 1PPS output.
- Use the GPS Receiver's status information (satellites and SNR) to make sure the temporary antenna installation is good and assure the best RF reception possible. Watch out for potential multipath reflections caused by tall buildings and reflective glass. Go to the roof-top, if necessary.
- Some stand-alone GPS Clocks may turn on their green Lock indication long before they achieve their best accuracy. Monitor the Altitude or Elevation to make sure it is no longer changing and its value makes sense for where you are at. We have found that this could be a good indication that the GPS survey has identified its accurate position and can now calculate time accurately.
- Minimize temperature variation as much as possible (also read the next point). No matter how good precision oscillators are, they all have some amount of temperature dependency. Since any frequency variation in the order of a few parts-per-trillion makes a big long-term difference in holdover error accumulation, try keeping the test equipment at the same temperature as the equipment room (e.g. 20°C/68°F) while doing the disciplining outside. For example, use the vehicle’s AC system during disciplining and the carrying case while transporting it inside (test set must be in sleep mode).
- In certain cases, it may be a good practice to allow normal/typical temperature changes during the disciplining process. This would allow the internal disciplining loop to build a temperature profile for the oscillator and use it later, in the holdover control process, to infer corrections if/when the ambient temperature changes. This functionality is usually available in integrated (fully-enclosed and tightly-packaged) precision oscillators, such as those using CSAC technologies with built-in disciplining control loop.
- Monitor the internal oscillator's Phase Alignment graph to get an idea of when it has reached an agreement with the GPS receiver and has become stable. A horizontal trend around 0ns would be a good indication that it may be safe to initiate the holdover. Keep in mind that this is internal indication is relative, since there may not be any other references to compare it to.
- Do not disable disciplining function to force the equipment into holdover. Instead disable or turn off the GPS receiver side. If no other options are available, disconnect the antenna. (Switching the disciplining function from ON to OFF is not recommended, as it may change some parameters in the oscillator and affect the corrections previously made by the disciplining process.)
- Make sure everything else required for the on-site test scenario is all set up and ready, so the test can be started as quickly as possible, after entering holdover. No time or accuracy is wasted waiting for other things.
- Consider leaving some margin for all the uncertainties, when measuring Absolute Time Error. For example, 100ns for the GPS receiver, 100ns for the disciplined clock and 100ns for the holdover drift. If the limit is ±1,100ns, perhaps the internal goal could be set to a ±800ns Pass/Fail threshold. (These numbers are presented for illustrative purposes only. This is not a specific recommendation.)
- Set reasonable test time limits. Not too short that TE or TIE traces could be misleading (a section of slow cyclic wander could be mistaken for offset) and not so long that too much holdover error accumulates.
- Consider that “extended holdover” may not be suitable for long-term testing, so look for the best practical holdover performance within 8 hours (work hours) or even 16 hours (overnight tests).
- At about 5 ns/m (1.5ns/ft), any significant cable-induced delays must be accounted for in the test setup, especially the antenna cable. If a local or portable reference (Cs or Rb) is used, make sure the cables carrying the reference signal and signal under test are of the same length, to cancel their respective delays.

- Learn about the oscillator technology used in the selected Frequency and Time reference. Identify and document its requirements: Warm up, temperature, overnight and long-term storage, handling (shock/vibration), calibration requirements and how to keep it in top performance. Then add "Equipment Handling Guidelines" to field test procedures.
- Visit and understand different typical sites, in which measurements will be performed, so test procedures align with the reality end users (field technicians) will be facing.
- Training the work force is very important. After all, phase could be something new to them. Understand experience levels and training needs for the field crews who will be performing these verification tests. Do they have frequency sync experience? Do they need phase/timing/time synchronization training?

Conclusion

When properly understood, and used correctly, GNSS (GPS) Disciplined Oscillators and Holdover are indeed powerful tools. They enable Phase and Frequency accuracy and stability measurements in places that would otherwise be impossible or cumbersome to test, due to the lack of local references. It many cases, it may be the only option available to verify time accuracy or one-way latency. End users should become familiar with this technology and understand the way it works, in order to use it properly, achieve the best performance and obtain valid results. Training is one of the most important aspects for verifying Time Error in the field.

Forced holdover offers limited test time, which can vary depending on the application, oscillators' quality, disciplining settings, temperature variations and local conditions. For practical purposes, the oscillator's corrected frequency accuracy in holdover mode should be much better than 0.01 ppb (10 ppt or 1.0E-11), in order to provide more than 10,000s (02:46:40) of transportation, setting and test time, with less than 100ns cumulative error. This does not imply that two hours would be the suggested limit. It should be much longer indeed. It all depends on the oscillator's accuracy and stability, and the uncertainty levels the application can tolerate.

Select test equipment carefully, taking into account the practical aspects of all your field applications. Oscillator technology (accuracy and stability), built-in GNSS, ease of use, battery operation, autonomy, size and weight are very important factors. Also consider if your field crew is specialized (synchronization oriented) or multi-purpose (test multiple services and technologies), so the instrument fulfils all their requirements. For example, base station installers may also require CPRI or OBSAI testing for the RAN/DAS part of the deployment.

Evaluate and compare instruments under actual use cases, in the field. Controlled lab tests may provide some insights, but true field tests will reveal important characteristics, advantages or limitations of each test sets.

Instruments and references used in remote sites should provide detailed status information about their GNSS receiver, oscillator and disciplining process. After all, that may be the only information field crews may have available to assess the instrument's readiness or decide when to initiate a successful holdover, if required.

Implement practical test procedures that fit the local needs and requirements, with end users in mind. Set practical limits and realistic expectations.

Training field crews is one of the most important aspects of this process.

"Wherever there is access to accurate, stable and traceable frequency and timing references, use them!"

Glossary of Useful Terms and Acronyms

1PPS	One Pulse-per-Second (usually aligned to the UTC second to distribute time around the world).
BIS	Bringing Into Service (final service verification and acceptance tests).
BNC	Unbalanced coaxial connector used for data and clock signal.
CPRI	Common Public Radio Interface (between cellular radio equipment controllers and radio equipment)
Cs	Cesium (Caesium). In synchronization, it refers to the atomic oscillator element.
CSAC	Chip-Scale Atomic Clock technology.
cTE	Constant Time Error (theoretical constant phase delay component created by cables and electronics)
DAS	Distributed Antenna System (cellular).
dTE	Dynamic Time Error (theoretical variable phase delay component due to phase noise, wander, jitter, PDV,...).
DOCXO	Double Oven Controlled Crystal Oscillator.
DUT	Device Under Test.
eNB	Evolved UTRA Node B or eNodeB (LTE base station).
GLONASS	Globalnaya Navigatsionnaya Sputnikovaya Sistema (Russian GNSS).
GNSS	Global Navigation Satellite Systems (generic term).
GPS	Global Positioning (satellite) System (USA's GNSS). It also refers to GPS time standard.
GPSDO	GPS Disciplined Oscillator.
I&M	Installation and Maintenance.
LTE	Long-Term Evolution (high-speed cellular technology, also known as 4G).
MTIE	Maximum Time Interval Error.
NTP	Network Timing Protocol.
OBSAI	Open Base Station Architecture Initiative (between cellular radio equipment controllers and radios).
OCXO	Oven Controlled Crystal Oscillator
OWD	One-Way Delay (latency)..
ppb	Parts Per Billion (1/1,000,000,000, 1x10^{-9} or 1E-9).
ppm	Parts Per Million (1/1,000,000, 1x10^{-6} or 1E-6).
PDV	Packet Delay Variation (Ethernet).
PTP	Precision Timing Protocol (IEEE 1588™).
PRC	Primary Reference Clock (Main frequency reference).
PRTC	Primary Reference Time Clock (Main Phase, timing and/or time reference).
RAN	Radio Access Network (cellular).
Rb	Rubidium. In synchronization, it refers to the atomic oscillator element.
RTD	Round-Trip Delay (latency).
SMA	Unbalanced coaxial connector (threaded) used for RF and clock signals.
SUT	System Under Test.
TAI	Temps Atomique International (International Atomic Time). TAI=GPS+19s
TDD	Time Division Duplexing (devices transmit and receive in the same frequency but at different time).
TC	Time Constant (time window used by disciplining control systems to assess and maintain sync).
TCXO	Temperature-Compensated Crystal Oscillator.
TDEV	Time Deviation (time stability of phase vs. observation interval τ of a clock source).
TE	(Absolute) Time Error (also known as Phase Error). Theoretically TE = cTE + dTE.
TIE	Time Interval Error.
ToD	Time of Day. Label with the yyyy/mm/dd hh:mm:ss information related to the last 1PPS pulse.
T&M	Test and Measurement (equipment or industry).
UMTS	Universal Mobile Telecommunications System.
UTC	Coordinated Universal Time (this is the primary commercial standard by which the world regulates clocks). UTC=TAI+LeapSeconds
UTRA	UMTS Terrestrial Radio Access (Node).
XO	Xtal (crystal) Oscillator (usually Quartz).

"Minimize ambient temperature changes during the disciplining, transport and extended holdover process."

Notes

About VeEX Inc.

Founded in 2006 by test and measurement industry veterans and strategically headquartered in the heart of Silicon Valley, VeEX Inc. provides innovative Test and Measurement solutions for next generation networks, services and communication equipment.

With a blend of advanced technologies and vast technical expertise, VeEX's products diligently address all stages of network design, verification, deployment, maintenance, field service turn-up, troubleshooting and integrate legacy and modern service verification features across DSL, Fiber Optics, WDM, CATV/DOCSIS, Mobile backhaul and fronthaul/DAS (CPRI/OBSAI), next generation Core & Transport Network, Fibre Channel SAN, Carrier & Metro Ethernet technologies and Synchronization.

The VeEX team brings simplicity to verifying tomorrow's networks.

VeEX
The Verification Experts