A Beginner's Guide to GPS Disciplined Oscillators

M. Greenman ZL1BPU
May 2007

Introduction

This paper has been prompted by the recent availablity to Radio Amateurs of a small number of GPS Disciplined Oscillators (GPSDO), and as a result, there have been questions from new owners about how they work, how to use them, what to look for, and how to ensure their performance.

This paper introduces the idea of GPS reference, points out the importance of *accuracy*, *stability*, *precision*, and *ageing*, describes how the GPSDO works, and how to determine how *well* it works, in terms of these parameters.

It is hoped that the information here can be kept as much as possible to layman's terms, and any technical terms that cannot be avoided will be *italicized* and defined.[1]

Advantages of GPSDO

The traditional frequency reference in the radio shack was for many years a simple crystal oscillator, designed to be reasonably constant in performance over operating position, temperature and time. The war-time BC221 and LM10 frequency meters contained such a device, while during the 1960s and 70s, it was common for receivers to include a 100kHz oscillator for use as a frequency marker. Such devices had accuracies measured in parts per thousand, and so small matters such as ageing, load, power supply and temperature dependence did not arise.

With the arrival of the digital age, frequency counters started to enter the ham shack, and with frequency resolution of about 1ppm (part per million), one would expect the reference therein to provide performance of this level as well, but typically this was not so! Even today, few counters in ham use have better than 1ppm performance, despite having 1Hz resolution to many MHz. Modern transceivers and receivers include frequency displays, but again, the performance rarely approaches 1ppm.

With the use of higher carrier frequencies, narrower-band transmissions, and new digital techniques, the requirements for stability and accuracy are also increasing. It is not uncommon now for Amateur techniques to demand 0.1ppm or better performance, and even 0.01ppm in some cases.

Lack of accuracy in a frequency counter or receiver doesn't particularly matter, if some other means of comparison is available. If a frequency counter reads 5ppm high, it can still be used for 1ppm measurements provided the resolution exceeds 1ppm and the size of the error is known.

The trouble comes in providing this necessary in-the-shack reference that can be depended on to be MUCH better than anything else that needs measurement. This reference also needs to be the same every time you use it. The GPSDO can do this, because it is (among other references hams can afford) unique in being automatically *traceable*.

---

[1] A simple introduction to GPSDO applications can be found at http://trl.trimble.com/docushare/dsweb/Get/Document-8439/

Accuracy, Stability and Precision

John Vig wrote a famous paper[2] on the subject, and introduced the concept of the target shooting target to illustrate the difference between these terms.

Precision is about putting all the bullets through the same spot on the target, not necessarily the bulls-eye – with precision, results are repeatable.

Accuracy is about hitting the centre of the target – the results are centred on the bulls-eye.

Stability is about getting the same results yesterday, tomorrow, day and night, so that at any time you know what to expect – every target has holes in the same place.

The same applies to oscillators. A *precise* oscillator gives the same measurement every time. An *accurate* oscillator is on the expected frequency. A *stable* oscillator changes very little from day to day. Of course all these figures of merit are relative – some oscillators are better than others in different respects, and no oscillator is perfect. Even the famous Caesium Standard (which is accurate by definition) is not perfect – because it's not as precise as other types.

In the design of very good quartz oscillators, these effects are well controlled and understood (for example good buffering and temperature control are used to ensure precision, frequent calibration to ensure accuracy). The remaining problem, stability, will be affected mostly by the *ageing* effects of the oscillator. Mostly this is due to physical changes in the quartz crystal, but other oscillator circuit components such as capacitors also age.

If a very good quartz oscillator is combined with a means of automatic long term calibration to correct for ageing, we have a reliable, self-calibrating frequency standard. Enter the GPSDO!

GPS and Traceability

The US Department of Defence Navstar system (commonly known as the Global Positioning System) consists of ground control stations and a constellation of orbiting satellites, each containing frequency standards, and each transmitting their precise time information.

These satellites are 'kept in line' with the DOD reference time and frequency standards, which are kept in line with the world time and frequency standards (UTC), i.e. are *traceable*. It is important to realize however, that the DOD reserves the right to fiddle with the traceability for strategic reasons.

Potentially the DOD could cut the connection to UTC at any time, and have in the past intentionally muddied the timing information (euphemistically called 'Selective Availability' or SA) in such a way that while slow measurements of time and position are possible, high speed measurements (such as are necessary for fast moving devices) are rendered inaccurate. At present SA is off. GPSDO devices are designed to use filters to remove the effects of SA, and are effective because slow adjustments are practical.

The signal from each satellite is measured in the GPS receiver, and the time of arrival (according to a stable local clock) is compared with the satellite's transmitted time. From this can be inferred the exact distance to the satellite. Since the satellite position is also known, this gives the beginnings of position information for the GPS receiver. By repeating these measurements and calculations for several satellites, the position of the receiver can be determined in latitude, longitude and altitude. The position is usually 'over determined', which means that although only three sets of data are required for a 'fix', six, eight or even up to 12 satellite measurements are made, and the ones which provide the best agreement are used to determine the position.

---

[2] *Introduction to Quartz Frequency Standards – Accuracy, Stability and Precision*, John R. Vig, http://www.ieee-uffc.org/freqcontrol/quartz/vig/vigtoc.htm, see page http://www.ieee-uffc.org/freqcontrol/quartz/vig/vigaccur.htm

The timing of signals arriving from the satellites is also subject to other sources of variation, which affect the accuracy of the computed position. These include ionospheric and atmospheric factors, signal reflections off buildings and large moving objects such as aircraft, and timing resolution limitations (quantization errors) in the receiver.

Time Reference

Much the same approach to navigation applies to determining the local time. The GPS receiver used for timing purposes (often called a 'timing engine') uses the time information from the satellites to over-determine the time reference and compensate for its own oscillator[3]. The GPS receiver generates a pulse every second (the time base on which measurements are made), and the timing of this pulse is corrected to the over-determined time solution, typically to a resolution of about 25ns (using a 40MHz processor).

GPS units designed specifically for timing applications generally include a special feature called 'T-RAIM', which is able to not only over-determine the time with high resolution, but to decide if (a) any particular data contributing to the average solution is significantly different in time (say outside 100ns from the mean), and so should be ignored, and (b) if the local one second reference cannot be guaranteed within a certain limit (typically 100ns). This determination may be because the GPS solution is poor, or because the local timing reference is still being adjusted. The result of these two features is that the GPS unit can provide a more reliable one second reference, or can warn when the reference is not reliable.

The one second signal is called the 'One Pulse per Second' (1pps) pulse, and this usually has the leading (positive going) edge within 100ns of the true UTC second event. The trailing edge of the pulse can be 20us to 10ms later, and is often programmable. The leading edge can also frequently be shifted slightly, to correct for delay in the receiver antenna cable, which can easily be 100ns in 20m of cable.

This attention to the 1pps pulse is fully justified, because this pulse is the basis for the operation of the GPSDO. While each 1pps pulse is known to only within (say) 100ns, so is the next pulse, and the one after. Second by second, this is an accuracy of one part in 10^7. After 1000 seconds, we still know the timing of the next pulse to within 100ns, but by now the time measurement over 1000 seconds is known to one part in 10^{10}. By continuing the process indefinitely, the accuracy with which traceable time is known easily exceeds the performance of any local reference.

---

[3] The GPS unit's oscillator is not corrected – it's affects are compensated for, in controlling the receiver synthesizer, the data sampling process, and the determination of time. It is important for a GPS receiver reference oscillator to be stable, in order to achieve and maintain its 'fix', but the oscillator need not be especially accurate.

Inside the GPSDO

The GPSDO uses this inherent long term accuracy of the 1pps signal to gently steer the local reference into perfect timing. Four main components make up the GPSDO – the GPS receiver, the reference oscillator, the time measurement system, and the control system.

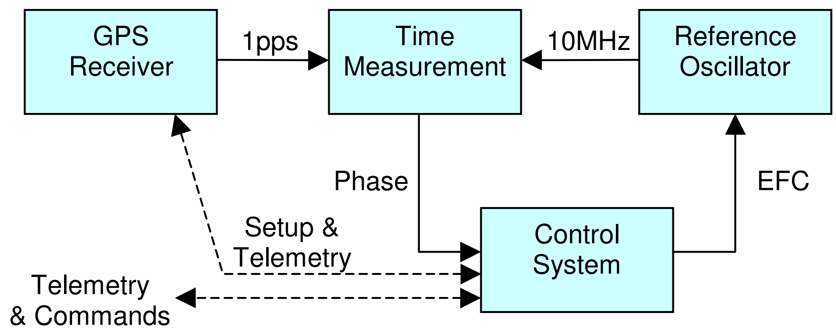

Fig. 1. Block diagram of a typical GPSDO

The GPS Receiver

Unlike an ordinary consumer-grade navigation GPS receiver, in the GPSDO the GPS receiver has no display, no ability to navigate or store waypoints, and generally no way to communicate with the outside world at all. All access to it is through the control system.

The GPS timing receiver is also not intended to operate while on the move, and the internal parameters and filters in the GPS receiver are designed to maximize performance while stationary.

However, the GPS timing receiver does have the ability to survey its own position very accurately – remember that with or without SA, the apparent position of the receiver tends to wander around due to inconsistencies in timing, both those in the receiver, and those caused by propagation from the satellites. By performing the site survey, the unit is able to determine its position with greater than normal accuracy, and that in turn allows it to then determine time with greater accuracy.

Unlike the ordinary hand-held GPS unit, the main output of the GPS timing receiver is the 1pps pulse. This pulse is also usually present (but can be wildly inaccurate) when the receiver has no good 'fix', and so it is important for the GPSDO system to know when the 1pps timing can be trusted. This is the responsibility of the T-RAIM feature and the GPSDO control system.

Once the survey has completed (initially this may take 24 hours, and on subsequent restarts may occupy an hour or less, or even be omitted if the position is stored), the unit switches into a 'position hold' mode. In this mode, the position is considered to be accurate, and all received position inconsistencies are used to improve the quality of the time determination. Since multiple satellites are measured at each determination, those best coinciding with each other can be used to maintain the accuracy of the 1pps signal. It is this 1pps signal that is used as the traceable reference.

It must also be obvious that the quality of the 'fix' is dependent on the number of satellites available. A GPSDO should always be operated with a very good, well positioned antenna.

The Reference Oscillator

The local reference in a GPSDO is typically a 10MHz oven oscillator. Quartz is an interesting piezo-electric material which has different mechanical, thermal and electrical properties in the three axes. In a reference oscillator the frequency versus temperature characteristics are very important. The effect of the different properties in two of the three axes can be illustrated by a graph of frequency versus temperature for an AT-cut crystal, where the angle of cut is varied slightly:

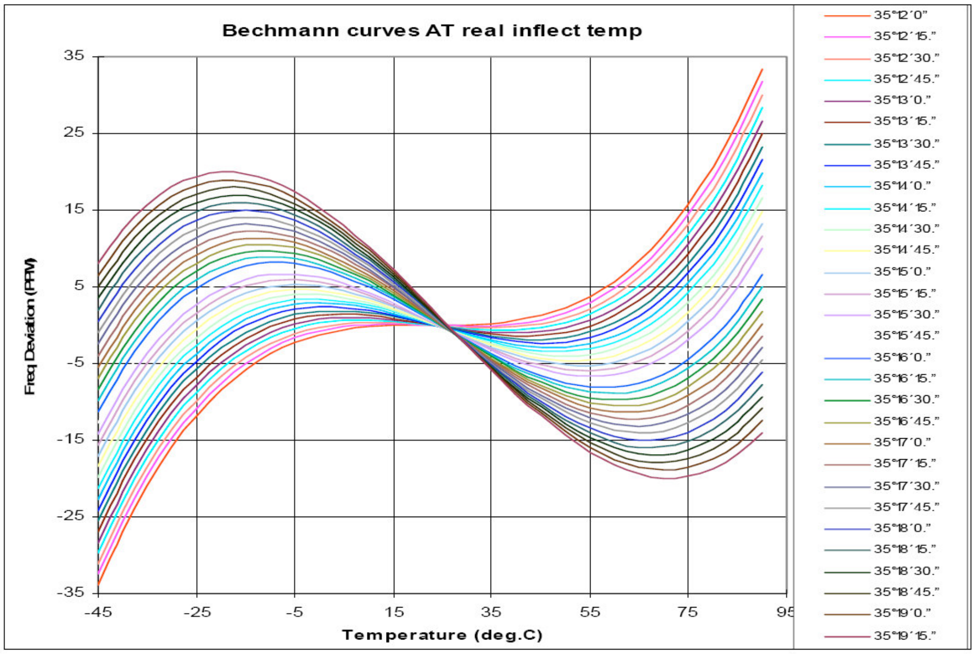

Fig. 2. Frequency vs Temperature for AT-cut crystals

Quartz crystals typically have a cubic-shaped frequency versus temperature characteristic, as illustrated clearly in the diagram, some with a plateau at around –20°C ±20°C, and a trough at around 80°C ±20°C.

In the best ovenized references, the crystal will be an SC-cut device, designed to have a frequency vs. temperature curve which has a substantially flat portion over several degrees. This is achieved through an adaptation of the AT-cut, by cutting the crystal at a specific Y- and Z- angle relative to the X-axis, in what is called a 'double rotation' cut.[4]

The crystal and its oscillator is mounted in a small oven which holds the crystal at the centre of the elevated temperature flat portion of its range (more convenient than cooling). In this way, the frequency of operation of the crystal is stabilized. Further, by careful crystal, electrical and mechanical design, and/or by adding a further oven around the outside (the double oven) the effects of not just different temperatures, but sudden changes of temperature, can also be considerably reduced.

The oscillator and crystal will have been designed for minimum phase noise, for minimum random variation over a longer time period (like low frequency phase noise), and also to have no jumps or dips in activity or operating frequency. It will have a control voltage input (typically called the Electronic Frequency Control or EFC), which allows it to be adjusted over perhaps 1 – 2 ppm about the nominal frequency.

[4] Most conventional frequency control cyystals are AT cut, where the axis of the blank is rotated in only one plane. SC-cut crystals are much more difficult to make accurately and repeatably, and are therefore much more expensive.

The oven of course includes heaters to maintain the stable temperature required by the crystal. The heater power is probably the largest contributor to the overall power consumption of the GPSDO, and is typically about 5W. The case of the oscillator is usually shiny to prevent radiation of heat and absorbtion of external heat, and the surface typically runs quite hot – too hot to touch. Lagging a good quality oven usually gives no advantage, but it can be beneficial to prevent draughts and changes of air flow around the oven.

Power supplies hum and noise, power supply voltage, load impedance, physical environment (for example protection from direct or intermittent heat, sunlight or draughts) and mechanical rigidity of mounting all can affect the performance of these oscillators.

If you had to go out and buy one of these oscillators, expect to pay from $200 to $1000!

The Time Measurement System

The GPSDO depends on a Time Interval (TI) counter to achieve its comparison between the 1pps signal and the local reference. This is in concept much the same as a conventional frequency counter, except that the counter is gated in a different way, the purpose being to measure the time difference between the 1pps pulse and the 10kHz reference – essentially a phase measurement.

Imagine a digital latch that is set by the arriving 1pps pulse from the GPS timing receiver, and reset by the 10MHz reference. The latch could be set for any time from close to zero, up to 100ns, the period of the 10MHz signal. Furthermore, the time the latch is set will vary with the phase (and amplitude!) of the 10MHz reference. Hence, if there was a way to digitally measure this time, and we could keep the 10MHz signal amplitude constant, we would have a digital number which represented the phase of the local reference.

A conventional digital counter can be used to recover the oscillator phase. What's more, the counter does not even need a very accurate timebase, because the 1pps pulse fulfills this role. See Figure 3. Imagine a digital counter with a 100MHz clock. At the 1pps pulse edge, the clock is allowed into the counter by a gate. This is controlled by the latch described above. When the next 10MHz reference comes along, the latch is cleared, the clock disconnected by the closing of the gate, the counter value is read, and the counter cleared ready for the next measurement (at the next second event).

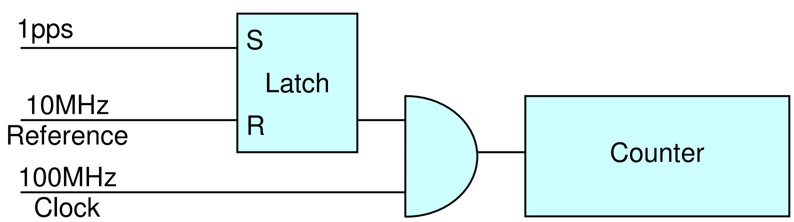

Fig. 3. A simple Time Interval Counter

The maths is simple. The biggest number that can be accumulated will be if the gate is open for 100ns (period at 10MHz), and in this time 10 pulses (at 100MHz) can be accumulated. Although the simple counter described has very little resolution (counts from 0 up to 10), it has quite enough resolution to form a useful phase measurement. Because the resolution is only 10ns, the 100MHz clock could easily have an error of 1% (0.1ns) and would still not affect results. This technique is called a TI Counter.

The phase measurement described is very coarse, and real GPSDOs use a slightly more complex system that gives greater phase measurement resolution. The reason for this is to achieve good resolution without compromising system frequency response. For example, if we were to divide the 10MHz reference down to 100kHz, with the same 100MHz clock and gating, we could have a new result every second with 1000 resolution steps rather than just 10. Instead of compromising the frequency response, this approach sacrifices absolute resolution (essentially each step represents 100 cycles of the reference), but in fact this isn't a bad thing since we are after long term performance. This technique is called a TI Average Counter, and is the basis for most high quality commercial Time Interval counters and some GPSDO control systems.

Other GPSDO systems do not use a TI counter, but instead use an equivalent analog sampling technique. The 10MHz signal is divided down to a convenient frequency such as 10kHz, and converted into a triangular waveform with very linear slope and accurate amplitude. This is like an analog equivalent of the incrementing numbers in the TI counter. The leading edge of the 1pps pulse is used to sample this ramp. The sampled voltage is held on an integrator (sample and hold acts like a switch and a large capacitor). The result is a voltage that represents the phase, and can be measured by a conventional digital voltmeter technique.

Control System

The purpose of the Control System is to adjust the Reference Oscillator to be exactly on the nominal frequency, essentially using a phase locked loop technique – adjusting the 10MHz oscillator phase so that it remains at a constant phase angle to the 1pps pulse, as measured by the Time Measurement System. This sounds simple, but there is a lot more to oscillator control than immediately appears.

The first point to be considered is that the Reference Oscillator has (by design) very good short term performance – very low noise, and because of isolation from the environment, no reason to change frequency except through very slow ageing. If we need to adjust it for any other reason than initial setting up and subsequent ageing, then we've not got a very good oscillator. Most oscillators of this quality can maintain one part in 10^{12} for an hour.

The GPS 1pps pulse does not have very good short term performance, and so to make use of it, the Control System must use long time periods, as previously explained. To complicate matters, in order to achieve reasonable performance in terms of settling time and control stability, the Control System must be able to make frequent control adjustments, best achieved if the time period is short! The answer is to use a long memory containing many timing samples (so time can be compared over a long period), and to have a new sample stored every second, thus maintaining a reasonably fast control system.

In a simple system, the phase measurement value might be used to directly control the oscillator EFC. However, a number of factors make this impractical. The gain would be low, the direction of control might be reversed, the centre-point of the control range might be inconvenient, and there would be no way to compensate for oscillator differences. Further, some sort of filter is required to average out the jitter in the GPS 1pps solution (remember the 1pps is only accurate over a long period), and with such a filter in place (invariably a software filter), a very sophisticated control system is required.

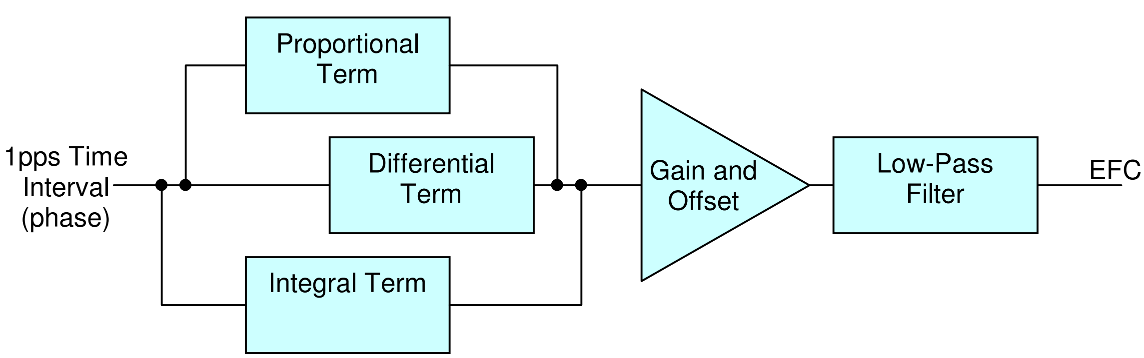

Fig. 4. A digital PID control system

Typically the control system applies a Proportional, Integral and Differential (PID) technique for control. Proportional feedback gives DC coupling to prevent drift. Integral control gives the system effectively infinite gain at low frequencies, while the Differential term allows the system to react quickly to step changes.

Since the control system is digital in nature, a digital filter is the most practical way to ensure that the reference oscillator is influenced only over the long term, to correct ageing, and not by short term GPS effects. Most systems include a clever type of low pass filter where the corner frequency reduces with time, allowing quick settling when power is first applied, and increasingly improved performance thereafter. Filter time constants typically start at 10 seconds and increase out to an hour or more.

Holdover

Most GPSDO devices encountered by Amateurs are surplus devices from telecom applications, where one of the most important attributes of the GPSDO is the ability to hold accurate timing when the GPS is unavailable[5]. A simple approach to this, and which is quite effective for short breaks, is to simply freeze the control voltage until the GPS signal is again stable, and the time measurement and control systems again have correct operation. Then the control system takes control again.

Unfortunately, aging and some other less well understood factors influence the holdover performance.[6] The longer the time without GPS correction, the more inaccurate the timing becomes – rather like the problem of 'dead reckoning' navigation. Without the occasional 'fix' on the sun with a conventional sextant and chronometer, the accuracy of the position of a ship calculated purely from course and speed becomes less reliable with time. The solution to this in nautical terms is to take account of currents in the sea, and drift off track caused by the wind. A very similar approach is taken in the GPSDO, using a system called a Kalmann Filter.

Without getting into the complexities of the Kalmann Filter, we can describe it as a learning system. A series of parameters are given to the filter as somewhere to start, and as time goes by, the filter recalculates these parameters and corrects them where necessary. Then, it looks at what has happened to these parameters over time, and when it has to carry on without GPS, it knows how much to make these parameters change until GPS is restored.

Yes, this sounds complicated, but as an example, look at just one term – perhaps the most important one – the one associated directly with ageing. The parameter of interest is the absolute EFC value. Let's say the value yesterday was 345678, and today it is 345679, and tomorrow we might reasonably assume it would be 345680. After a few days, we will know quite accurately how much the EFC is changing per day (in this case +1 step/day). So, if we suddenly lose GPS at any point, the Kalmann Filter can apply a value of 34567 (or whatever it was before the GPS was lost), and then adjust it by +1/86400 every second that GPS is missing. When GPS returns, the Reference Oscillator will be closer to where it should be than if the value was simply held constant.

So there are two reasons why setting up a GPSDO for the first time is a slow process –

- The GPS unit must find the satellites, determine the position, and then survey the site to give best accuracy to the position.
- The control system must lock and stabilize the phase of the oscillator, and give the Kalmann Filter time (several days) to learn the necessary parameters.

As a consequence, when you set up a GPSDO, you must expect to power it continuously for ever after – and seriously consider using a battery backup power supply.

---

[5] A technique known as 'holdover'. A typical telecom specification would require system timing to stay within 7μs over 24 hours.
[6] That's assuming the oscillator is unaffected by temperature, load and supply voltage!

System Telemetry

Although GPSDO units are designed to be fully automatic and self contained, the Amateur can learn quite a lot about the quality of the system and ongoing performance of the oscillator by looking at some crucial information relating to the system. Remember that the input to the control system is the 1PPS TI (oscillator phase measurement) and the output is the EFC (control voltage).

1pps TI

In the short term, the 1pps measurement is a good indication of the GPS timing receiver performance. There will always be jitter on the timing, and it is important realize that this is NOT jitter on the local reference, because in general it will not be (unless the unit is a cheap and simple 'GPS steered oscillator').[7] The jitter is as a result of GPS signal timing effects, GPS receiver timing resolution, and TI counter resolution limitations.

If the 1pps timing has significant noise, look carefully at the location and performance of the GPS antenna. If there are sudden jumps, look at causes in the Reference Oscillator. You can easily see the effect of a sudden jump by turning your GPSDO unit upside down – nothing should fall out, but the heat flow will be reversed, and you should see a transient which takes perhaps several hours to settle.

If there are transients that are not related to manually induced effects (Did you bump the unit?, Did the sun come in the window and warm it up? Is your power supply unreliable?), it is possible that they are caused by micro-jumps in the reference oscillator crystal. Even so, the effect on the actual 10MHz reference will be quite small, and often not detectable by conventional (i.e. radio reception) means. The highest quality oscillators have fewer jumps, usually as a result of a combination of scrupulous cleanliness during manufacture, a careful testing and selection process during manufacture, and great care taken in the design of the crystal.

The long term (average) value of the 1pps TI should be constant. A shift of 100ns over 24 hours represents a frequency error of about 10^{-12}. Commercial GPSDOs intended for telecoms applications use proportional control, and adjust the reference phase to be close to zero, and this value should stay constant. Other units may also hold the phase constant, but at an arbitrary angle.

The ZL1BPU VNG-in-a-Box has proportional control, but because the TI measurement is made at a multiple of 1Hz, it holds at one of several exact phase angles. Simpler units (such as the ZL1BPU GPSClock) do not have proportional control, and the 10MHz phase can be stable at some arbitrary value, or even change slowly.

Most commercial units provide commands whereby you can interrogate the instantaneous 1pps TI, but it is plotting the TI on a graph over hours or days that is so useful.

EFC

The Electronic Frequency Control is the voltage returned to the reference oscillator to keep it going with an accurate reference phase angle. Voltage controlled OCXO (V-OCXO) devices have a control range of 0 to +5V, -5V to +5V, 0V to 10V etc, and you could monitor this with a voltmeter. However, you would also induce noise and degrade the performance.

Fortunately the GPSDO control system generates the EFC as a digital value, which is subsequently converted to an analog voltage to control the oscillator, and this digital value can be interrogated. This value can be a 16 to 24-bit number, and can be an absolute value (say 0 to 65535) or a signed value (say –32768 to +32767).

---

[7] Where a cheap and not very stable room temperature oscillator is forced to track the GPS 1pps, not the other way around!

The relationship between the EFC value and the control range of the oscillator is determined by the D-A converter gain and the voltage versus frequency characteristic of the OCXO.

The important numbers to know are the minimum, the maximum, and the centre EFC values, because the operation of the oscillator can be assessed from these values and the current working value.

Ideally, and when new, the OCXO will operate close to the nominal frequency with the EFC at the centre of the range. As the device ages, the operating point will slowly shift in one direction, and it is a good idea to note this value every few months in order to determine how long the unit will stay operating correctly before some type of recalibration is necessary. Hopefully this will exceed the useful life of the device (or the owner!).

Once the EFC value reaches about 90% or has dropped to 10% of the available range, the ability to correct for transient errors becomes impaired, and the unit may not recover from a transient, or may not settle at all after power failure. Most commercial GPSDOs will bring up an alarm when theEFC reaches an arbitrary limit. Well designed units should have enough control range to operate for more than 10 years without recalibration.

The process of recalibration is NOT for beginners. It involves setting the OCXO to the centre of its range and adjusting the oscillator mechanically or by component substitution until the OCXO again operates at close to the nominal frequency. Generally this must be done independently of the GPSDO, although the ZL1BPU VNG-in-a-Box and GPSClock do provide a calibration mode where the OCXO is held at mid-range and the phase can be plotted and the oscillator adjusted for slowest change of phase.

Many OCXOs are welded or soldered closed, and are sealed in a hermetic environment filled with dry Nitrogen.[8] Any attempt to open these to make internal adjustments will inevitably impair the subsequent performance. This type of oscillator is essentially calibrated for life.

Other Parameters

It is generally useful (in order to assess ongoing quality of performance) to know how many satellites the GPS receiver is using to determine the fix, and how many satellites it is tracking. The signal strength (C/N or Carrier to Noise ratio) of each of the satellites can be expressed graphically, as a histogram. A satellite constellation map (satellite position, Azimuth vs Elevation) can indicate which parts of the sky the receiver can see well.

GPS time and date are useful in a general way, but care needs to be taken to understand the time value provided by the GPSDO. The GPS satellite system keeps time in UT1, which is the time since a given arbitrary point, assuming 86,400 seconds per day, and the usual arrangement of days per month and leap years. We frequently call this time ‘GPS time’.

Civil time (and Universal Coordinated Time, UTC) is based on an average of solar time. Since the rotation of the earth is slowing down slightly, and therefore gets behind UT1, the time observed from the GPSDO may (or may not) be different to UTC.

Every now and again an extra second is added to the UTC time in order to keep it within one second of solar time. This is done roughly once per year, either on 30 June or 31 December, at midnight, where the last minute of the day is given an extra second (called a leap second). At present, UTC is running 14 seconds behind.

The time kept in the GPSDO is usually GPS time. The current UTC offset is always broadcast from the satellites, along with warning information about the next leap second event. Since the GPSDO unit knows both GPS time and the current UTC offset, it can sometimes be set to provide UTC rather than GPS time. Be aware though, that not all time products from the unit may be changed in this way.

---

[8] Oscillator components, including the crystal, are adversely affected by humidity.

How do you tell which time setting you have? Simple – listen to the local (civil time) time pips on the hour from a national radio network. If your GPSDO time is 14 seconds ahead, you have GPS time. If it is correct, you have UTC. If it is late, you have delays in the computer! International time signals on HF (WWV etc) and time codes on LF (controlling radio clocks) also provide UTC.

Civil time is also used to control clocks in CDMA cellular phones and many computer networks, although since seconds are not displayed, you will need to watch closely to see the minute tick over. UTC is also available from any time servers around the world, but be aware of network propagation delays.