A Different Approach to Control of a Rubidium–based GNS Disciplined Frequency Standard

by Precise Time and Frequency, Inc.

Many applications utilizing precision frequency references use an OCXO as the local oscillator to provide improved accuracy and stability, however it is not uncommon to require the additional stability and holdover performance provided by incorporating a rubidium atomic clock as the local oscillator in a GNS disciplined standard.

Many of these systems also require a precision timing signal (e.g. one pulse per second or 1PPS) to synchronize in the time domain to UTC (Universal Time Coordinated).

Disciplining (or "steering") either an OCXO or a rubidium oscillator to UTC while at the same time preserving the inherent excellent stability performance of the rubidium is not a trivial task. Recent development work on a new product at PTF yielded the opportunity to consider a different strategy for control implementation.

This paper describes both the traditional OCXO control and an alternative method to the more commonly used PID type control loop used for the rubidium.

Disciplining to 1PPS, Control Loop Basics

In a conventional GNS disciplining method, there will be an incoming reference signal, usually a 1PPS pulse from a GNS engine, and this is compared to an output 1PPS pulse, usually derived from a low noise oscillator (typically an OCXO) running at a frequency of 10 MHz. The resulting phase measurement represents the timing difference between the two 1PPS signals, and is used to drive the frequency of the oscillator in such a way as to reduce the phase error between the two 1PPS pulses to zero to achieve the required frequency accuracy and stability.

Typically, the target specification for modern day frequency standards is to achieve a frequency accuracy of better than 1E-12, averaged over a period of 24 hours with a stability (expressed in terms of Alan deviation) of 1E-11 or better at 1 second, reducing to better than 1E-12 at 10,000 seconds and longer averaging times, with a 1PPS synchronization accuracy, expressed as a 1 sigma variation, of 20ns or better.

In concept this is a fairly straightforward, however several factors make the implementation somewhat challenging:

1. The input 1PPS pulse is produced by decoding the transmitted signals from one or more satellites (GPS, Glonass, etc.) and by the time these signals reach the Earth, they are not only very small, typically -130 dBm or less, but they have also been subjected to several noise inducing factors, such as ionospheric variations. This gives a variable transit time from the transmitter on the satellite to the receiver on the ground. The result is that the reference 1PPS signal can be quite "noisy," with variations of 100 nanoseconds or more.

2. During "warm up," the phase error between the input 1PPS and the 10 MHz derived 1PPS can be substantial. If steps are not taken to quickly reduce this error to a

manageable amount, the time to synch the two signals would be unacceptable (possibly many hours).

3. There are conflicting requirements between obtaining an accurate and stable 10 MHz frequency reference as quickly as possible, while at the same time eliminating the output 1PPS offset by adjusting the 10 MHz reference frequency.

To reduce the impact of factor 1 (inherent noise in the GNS signal), a technique of input filtering is employed to reduce the noise of the signal going into our control equation. In addition, the objective is to average the input signal over as long a period as possible to again reduce the inherent noise. In practice, the maximum averaging time for an OCXO based reference is around 100 seconds, before the aging of the OCXO becomes significant, whereas for a rubidium standard the averaging time can extend out to 0.5 days or longer; see the explanation below.

The impact of factor 2 is almost in contradiction to factor 1. In order to deal with rapidly changing frequency/phase offsets during a warm up period, it is necessary to control the local oscillator in a shorter time frame, thus minimizing the accumulated offsets.

Finally, for factor 3, the slower and smaller the changes applied, the better will be the accuracy/stability characteristic (up until reaching the limit imposed by local oscillator aging).

The end result of these considerations is that the ideal solution to providing an optimized performance is to have a control loop that adapts to the different situations and provides the necessary control accordingly. In practice, it has been found that this is best implemented with a 3 stage control, going from a fast loop (at warm up), and using a medium loop to transition from warm up, to a slow loop for steady state operation.

Quantifying the parameters is dependent upon the specific characteristics of the components used, but as a guide, the derivation of typical values is shown below:

OCXO Parameters

Typical OCXO aging = 1 E -9/day

therefore in 100 seconds

typical aging = (1 E -9 / 86,400) x 100 = 1 E - 9 / 864 = 1.1574 E-12

If the control loop time constant is too long, the oscillator aging will take the frequency outside of the target range of <1E-12.

A typical OCXO will have a control voltage input allowing adjustment of the oscillator frequency over a limited range, typically allowing offsets of around +/- 1E-7 for an input voltage of +/- 5V.

A frequency change of 1E-7 on 10 MHz represents 1Hz. In other words, the maximum electrical tuning range of the oscillator is +/- 1Hz.

If a 24 bit digital to analog converter is used to drive the control voltage each step represents:

1Hz/2^23=1/8,388,608 = 0.00000011920928955078125 ~ 1.2 E-7 Hz

this represents a (minimum) "fractional frequency" offset change (df/F), or resolution, of approximately

1.2 E-7/1E7 = 1.2 E-14

These numbers are important, given the target for most precision frequency references to provide a frequency accuracy of better than 1E-12 over a 24 hour period. The resolution of the control at around 100 times the required accuracy therefore gives a good chance of achieving the requirement.

Rubidium

Most rubidium standards are designed as stand-alone precision frequency references, many with the built-in capability to be disciplined to an external one pulse per second (1PPS) pulse typically derived from a GNS or other reference.

While this is ideal for straightforward frequency reference applications, a somewhat more sophisticated approach is necessary if the rubidium module is to be integrated into a more feature rich frequency and time standard.

In this case, it is more convenient and flexible to externally derive and control the output 1PPS signal from the generated 10 MHz reference provided by the rubidium module. During the warm up phase, the rubidium output frequency may have a relatively high fractional frequency error (5E-8 during the first few minutes). This offset translates into a phase change of around 50 nanoseconds per second, which quickly adds up.

Once initial lock has been achieved however, after approximately 5 minutes, the rubidium output frequency will typically be within 5E-10, or a phase change rate of around 0.5 nanoseconds per second. At this point therefore, most control implementations include a technique whereby the 1PPS signal is "jam synched"*see note 1 to align the rubidium derived 1PPS with the external (e.g. GPS) 1PPS.

Rubidium Control Philosophy

The different approach to control of a rubidium standard outlined in this paper is intended to take maximum advantage of the extensive development effort already expended on producing a highly stable stand–alone atomic frequency reference, while implementing the additional requirements demanded by a precision time and frequency reference standard disciplined to an even more accurate (in the long term—several days/weeks) reference source such as GPS.

Once again, it is convenient to utilize a control loop that adapts to the rate of change of the frequency offset, implemented in three basic stages; fast, medium, and slow.

In this case however, the control is implemented via a digital (RS232) command, instead of an analog voltage level. Most rubidium standards provide for a digital control input. The module used in this discussion allowing for incremental control changes to the frequency offset (i.e. resolution) of the order of 5E-13.

This is very adequate for the initial fast control loop, where the objective is to control the rubidium frequency offset to within around 5E-12 while maintaining the 1PPS phase error to around 200 nanoseconds or less. After a warm up period, usually less than 30 minutes, the rubidium output frequency stabilizes, reducing the amplitude necessary control adjustments.

For those that are interested, this control can be implemented as a classic Proportional–Integra–Derivative (PID) control with a loop bandwidth of around 2 milli Hz, i.e. time constant of 1/1E-3 = 1000 seconds.

The control equation is shown below:

CV = (k1 x Current Phase Error) - (k2 x last Phase Error) + (k3 x last CV);

Where:

CV = control value

Phase Error = Error in nanoseconds between 1PPS derived from the rubidium 10 MHz and the external 1PPS

k1, k2, and k3 are the difference equation coefficients

where τ is the sample time :

$$k1 = \frac{(2\,\tau_z + 1)\cdot\alpha}{(2\,\tau_p + 1)}$$

$$k2 = \frac{(2\,\tau_z - 1)\cdot\alpha}{(2\,\tau_p + 1)}$$

$$k3 = \frac{2\cdot\tau_p - 1}{2\cdot\tau_p + 1}$$

Once the rubidium module has stabilized, the frequency offset of the output remains very constant for long periods of time, with a typical aging rate of around 1.7 E-12 per day. The objective of the control loop at this point is twofold;

1. Adjust the frequency offset to bring the frequency to within <1E-12

2. Minimize the phase error between the 10 MHz d erived 1PPS and UTC

To obtain the best possible performance from the rubidium module, it is advisable to make the fewest adjustments necessary to satisfy the above two requirements as the long term target frequency is less than the inherent short term noise in the process. In order to achieve this, a different strategy is employed for the medium and slow control loops.

Rather than adjusting the module on a second by second basis, the 1PPS error (between 10 MHz derived PPS and GPS derived PPS) is measured and stored each second. At the end of a time period which

Precise Time and Frequency, Inc.
Con't from pg XX

is dependent on whether the medium or slow loop is being employed, a linear regression is applied to the data to calculate what the average slope of the frequency offset has been over that period.

In addition, the phase error is calculated, and a figure is calculated for the additional offset change needed to eliminate the phase error over the next sampling period.

These two calculations are then combined to provide the level of frequency adjustment needed. An analysis is shown in the following section.

The linear regression relates to two elements, namely the "slope" and the "y intercept" of the resultant fit. These two elements relate to the rate of change of the frequency offset and the phase error, respectively. The two parameters can be determined mathematically as follows:

Slope $$m = \frac{n\sum xy - \sum x \sum y}{n\sum x^2 - (\sum x)^2}$$

y intercept $$b = \frac{\sum x^2 \sum y - \sum x \sum xy}{n\sum x^2 - (\sum x)^2}$$

where:

n = total number of samples i.e. total number of seconds over which measurement is taken

x = number of the individual sample

y = value of the sample i.e., phase error of the individual sample

so m represents the rate of change of the frequency offset and y represents the phase offset.

Now to consider the phase offset; phase offset is represented by:

phase offset = $\frac{\sum_1^n y}{n} + \frac{m\,n}{2}$

average phase offset + 0.5 x increase in phase offset

Our two values to compensate then are the rate of change of frequency offset, represented by m, and the absolute phase offset shown above.

In terms of control values, these are calculated by:

$CV_{df/F} = m / r$

and

$CV_{phase\ err}$ = phase error / r x n

where:

m = slope (calculated above)

r = resolution (change in frequency for minimum applied change)

n = number of seconds over which the correction is to be applied

and our total control value is:

$CV_{tot} = CV_{df/F} + CV_{phase\ err}$

The advantages of this approach are:

1. By applying the change over an extended period of time, the control resolution is effectively increased by a factor proportional to the time period

2. The good short term stability of the rubidium is not perturbed, resulting in excellent short term

3. Absolute phase offsets are taken out very slowly over a long period of time

One other aspect to consider is that in order to avoid instability or overshoot while trying to take out 100% of the error in one go, a small degree of damping is implemented on the control value before being applied.

It is important to remember that even a very small frequency offset of <1 E-12 can result in large accumulated phase errors:

For df/F = 1E-12, the rate of phase error accumulation = 0.001 nanoseconds/second

Therefore in one day (86,400 seconds), total accumulated phase error is given by:

Phase error (nanoseconds) = 0.001 x 86,400 = 86.4 nanoseconds

If nothing were done about this accumulated error, over one week it would start to become significant at around 0.5 micro seconds.

In practice, the approach described has proven extremely effective and can deliver frequency offsets of <1 E -12 while maintaining phase error with respect to UTC of around 20 nanoseconds.

For further information on this or other time and frequency applications, please contact Precise Time and Frequency, Inc. or visit the web site at www.ptfinc.com.

*Note 1 "Jam Synching" is a technique whereby a large phase offset (say 5 milli seconds) on the output 1PPS signal is removed by forcing it to be synchronous with a reference 1PPS. Typically the 1PPS is derived by dividing down a 10 MHz RF input (by 10,000,000) by means of a counter/divider.

The jam synch resets the count to 0 on the incoming reference 1PPS with the result that the derived 1PPS should fall within the range of +/- 100 nano seconds (the period of 1 cycle) of the incoming reference 1PPS . Higher precision can be achieved if a 100 MHz clock (divided by 100,000,000) is used as the source of the 1PPS, giving a resolution of 10 nanoseconds.

• PRECISION TIME & FREQUENCY, INC •