

Environmental Effects and Control Systems for GPS-Disciplined Clocks (GPSDC)

Andrew Novick

Time and Frequency Division
National Institute of Standards and Technology
Boulder, CO. USA
novick@nist.gov

John Clark and Demetrios Matsakis

Masterclock, Inc.
St. Charles, MO. USA
jclark@masterclock.com
dmatsakis@masterclock.com

Summary—GPS-Disciplined Clocks (GPSDCs) are sensitive to their environment in many ways, and the susceptibility varies with different GPSDC designs. In the short term, temperature shocks can perturb the frequency of the disciplined local oscillator, causing shifts in the time output. Assuming the GPS receiver module is not separately impacted, GPSDC will return to its steady-state time and frequency setpoint, at least until the next environmental shock. We have tested three commercially available GPSDCs with varying results. Performance of these devices in various conditions will be described along with mitigation strategies, including variance of the control loop parameters.

Keywords—time; GPS; metrology; GPSDC; temperature.

I. INTRODUCTION

GPS-Disciplined Clock (GPSDC) environmental sensitivity can be a serious problem for equipment operated outside a stable laboratory environment. In some situations, thermal equilibrium may not be attainable because the units are continually exposed to new temperature shocks. Some GPSDCs are much more sensitive to these effects than others. The more sophisticated units, in general, provide better thermal isolation for their oscillators and key electronic components. However, many GPSDCs follow the general design of a local oscillator disciplined to GPS via proportional-integral-derivative (PID) or phase-locked loop (PLL) steering. The magnitude of the associated PID gain control terms determines their response times to shocks of all types [1]. GPSDCs differ slightly from the systems described in [1] in that a GPSDC tunes a Voltage-Controlled Oscillator (VCO) whose frequency offset is determined by a directly applied voltage instead of by a digital input that determines the change in the oscillator's frequency from its previous value. Mathematically, the difference is that the derivative gains used in [1] are larger than those in a PID/PLL by 1. The devices under test (DUTs) used in this work can be better described as PI loops because they set their derivative gain to zero, which is equivalent to setting $g_d = 1$ in the formulas of [1] so that the time constants can be found as solutions to

$$0 = r^2 + (-2 + g_I + g_P)r + (1 - g_P) \quad (1)$$

Where $r = e^{-\tau/T}$, g_P is the gain associated with the offset in phase (also called P-gain), g_I is the gain associated with the integral of the phases (also called I-gain), τ is the interval between epochs (1 second herein) and T is the time constant in seconds.

It then follows that

$$e^{-\tau/T} = \frac{2 - g_I - g_P \pm \sqrt{g_I^2 + g_P^2 - 4g_I + 2g_I g_P}}{2} \quad (2)$$

These equations complement Laplace and Z-transforms [2,3], all of which make some assumptions that are not necessarily valid for all GPSDCs, such as that the VCO operates in a linear range, along with all the associated circuitry and filters. Given these assumptions, the criteria for stability ($T > 0$) are that $g_I + g_P < 2$ and $g_P < 1$. Figs. (1) through (3) show how different values of gains g_I and g_P can lead to decaying oscillations of various magnitudes. The common theme is that larger values of P or I lead to tighter control, but at the possible expense of overshooting the setpoint.




Fig. 1. Response of PI loop to a 10-ns/s frequency disturbance in several cases where the P-gain Pg_P equals the I-gain Pg_I . Tighter responses. For example the blue, red, and green curves correspond to (P,I) true gain values of 0.1, 0.2, and 0.5 respectively.




Fig. 2. Response of PI loop to a 10 ns/s frequency disturbance when the I-gain equals 0.5, for several values of the P-gain. The larger P gain values had tighter responses. For example the blue, red, and green curves correspond to P gain values of 0.5, 0.1, and 0.2 respectively.




Fig. 3. Response of PI loop to a 10 ns/s frequency disturbance when the P-gain equals 0.5, for several values of the I-gain. The larger I gain values had a tighter response. For example the blue, red, and green curves correspond to (g_p and g_i) values 0.1, 0.25, and 0.5 respectively

II. EXPERIMENTAL DESIGN AND DUT CHARACTERISTICS

Three single-frequency GPSDCs, two with a single oven-controlled crystal oscillator (OCXO) and one utilizing a double-oven-controlled crystal oscillator (DOCXO) were placed inside a programmable thermal chamber (Fig. 4). Of the OCXO-equipped DUTs, DUT1 displayed considerably less jitter. All three DUTs were fed signals from a common antenna, as was a fourth DUT that was kept in a normal laboratory environment. It served as a control and verified that there was no unusual behavior related to GPS as seen through the antenna and its antenna splitter port. Atmospheric effects and multipath effects were therefore largely identical for the DUTs [4, 5].



Fig. 4. The programmable thermal chamber

The pulse-per-second (PPS) outputs of the DUTs were measured against UTC(NIST) with individual time interval counters, at a rate of one measurement per second. These DUTs had features that allowed the option to model temperature and aging effects.

To estimate the control effects of the PI loops and internal learned temperature compensation (tempcomp), the three DUTs ran in holdover mode (without GPS input) as the oven temperature was varied over setpoints from +5 to +45 °C. An independent wireless sensor inside the chamber recorded temperature every five minutes. It could take up to an hour for the chamber to fully respond to a programmed temperature swing of 40 °C and up to 15 minutes for smaller swings. Device diagnostic measurements of the two OCXO-equipped DUTs' internal temperatures showed that they closely followed the external temperature fluctuations in time, but were internally hotter by 15 to 25 °C. The tempcomp adjustment comes from internal monitoring of the OCXO current, temperature measurement and electronic frequency control. External temperature changes could affect the GNSS receiver module and all the electronics in addition to the oscillator, however, the internal tempcomp adjustment of the DOCXO was much smaller, indicating that the most significant variation is due to the oscillator.

Another consideration is aging. To study this, the DUTs were subjected to temperature variations while in holdover mode (Fig. 5). The 900 s averages (center lines in the jitter) show if there is a slope in the frequency offset during the varying periods in between temperature changes. Interestingly, at low temperatures, DUTs 1 and 2 showed possible aging of 3 ns/s/day and 1 ns/s/day, while at 25 °C and 45 °C the rates were considerably less and not statistically significant. The aging of the frequency offset of the DOCXO-equipped DUT showed no significant temperature dependence. The small frequency slopes seen could possibly have a contribution from the DUT attempting to steer based on what would be expected from the last GPS observation, however this was ruled out due to the long length of the holdover period relative to the observed decay times when GPS was available.




Fig. 5. DUT frequency variations over short periods. This is during holdover conditions (no steering to GPS).

By converting the PPS measurements to frequency, a measurable and not unexpected tempcomp nonlinearity could be seen in all three DUTs, although the DOCXO-equipped DUT's tempcomp was significantly less. Fig. 6 plots the frequency expressed as first differences in ns, corresponding to a frequency offset in ns/s or 10^{-9} . In all plots the convention will be followed that the OCXO-equipped devices (DUT1 and DUT2) are shown as blue and red, and the DOCXO-equipped device (DUT3) is shown in green (horizontal line). Timing data are presented as UTC(NIST) minus each DUT.




Fig. 6. 12-hour averaged first differences (corresponding to frequency offsets of 1 ns/s) when DUTs were in holdover mode.

While Fig. 6 demonstrates an overall linearity, which the DUTs can be configured to learn through observation, there are minor variations in the OCXO-equipped DUTs when the slope is removed, as is shown in Fig. 7.




Fig. 7. Dependence of averaged first-differences (DUT frequency offsets) on temperature after removing the overall slope.

III. DATA AND ANALYSIS

Initial experiments in [5] that subjected the OCXO-equipped DUTs to rapid temperature shocks showed the reaction of the PPS output. Figs. 8a and 8b show high/low temperature shocks and also the temperature increasing in steps, and the resulting responses in the time offsets of the PPS output. The offsets were steered out by the DUT control loop, eventually reaching the original time offset (setpoint). The user-set gain control values (identified as P and I in the figures) are related to the gains g_P and g_I in equations (1) and (2), however the manual does not describe them in these terms. The factory default values for the DUTs are (P,I)=(4,25). It is assumed that doubling these gain control values doubles the gains. Fig. 9 shows that doubled values for the P and I gain terms led to faster responses and lower excursions of the time offset.

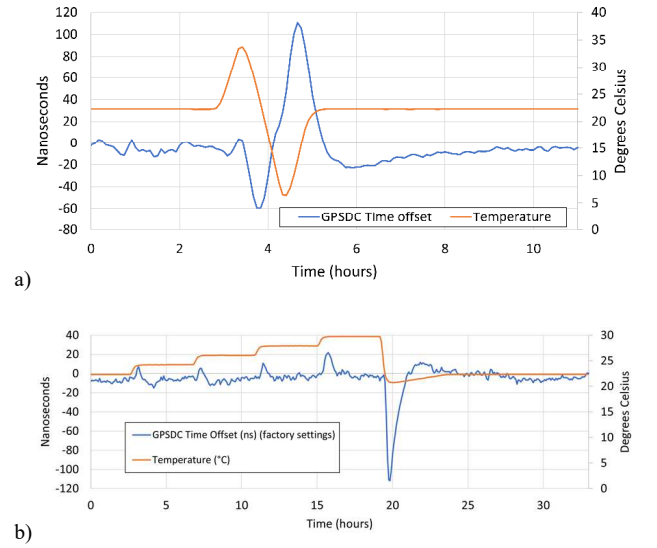


Fig. 8. Temperature shocks (orange) and DUT's PPS output (blue) referenced to UTC(NIST) using factory PID settings, with temperature compensation enabled.




Fig. 9. Temperature shocks (orange) and DUT's PPS outputs referenced to UTC(NIST) with factory settings (blue) and doubling of P and I controls (red, smaller values), with temperature compensation enabled.

To systematically evaluate the effect of changing the P and I gain terms on the DUTs, a series of experiments was run. Usually, DUTs were run without internal temperature compensation or aging corrections, the latter having little effect at the duration of the data runs. A typical cycle for the data

shown hereafter starts with the chamber temperature being held to nominal 25 °C, then raised 3 °C, returned to nominal, lowered by 3 °C, and then returned to nominal. The time at each phase was usually between two and four hours.

Initially, temperature variations were set to ± 20 °C, but it was observed that the response to the resulting extreme variations in the oscillator phase and frequency induced simple jumps towards zero that were usually multiples of 100 ns (Fig. 10).




Fig. 10. Phase variations during a 20 °C temperature change.

Fig. 11 shows the resulting curve with multiples of 25 or 100 ns removed from the data in Fig. 10. However, even after the $N \times 25$ -ns corrections were applied, some jumps of around 20 ns remained. This is clearly due to the instruments attempting to quickly get to their desired goal. While usually a good thing, subsequent tests used only 3 °C shocks to better follow the assumptions of the theory.






Fig. 11. Same raw data as in Fig. 10 after removing jumps that were multiples of 25 or 100 ns.

Figs. 12a and 12b show the phase and frequency response of the three DUTs when the gain controls terms for DUT1 and DUT2 were set to their factory default values of $(P,I)=(4,25)$. The outputs of these devices returned to their setpoints over several hours after a temperature change. As expected, DUT3 has a very small sensitivity to temperature variations due to its double oven. The tempcomp setting had been turned off to isolate the effects of the P and I controls for these tests. However, around MJD 60061.6, data logs showed that the tempcomp setting came back on for DUT1 (blue line) and the

difference can be seen in the latter portions of Figs. 12 and 13, where there was a lesser reaction to the temperature change.



a)



b)

Fig. 12. a) Phase and b) frequency response of the DUTs undergoing 3 °C temperature shocks. Temperature compensation for DUT1 was turned on beginning around MJD 60061.66 and the reaction to temperature change is lower than for DUT2.

Multiplying the P and I gain control values by four resulted in DUT2 having smaller excursions and faster responses to long durations of different temperatures, as shown in Fig. 13. DUT1 showed minimal effects due to the tempcomp setting being on, but they each had a phase excursion when the temperature changed by 3 °C and then changed back more rapidly (MJD ~60062.75).




Fig. 13. Phase response of DUTs undergoing 3 °C temperature shocks when (P,I) gain controls set to $(16,100)$. Temperature compensation for on during this period.

Setting the gain controls to even higher values (Fig. 14) resulted in oscillatory behavior, which would correspond to imaginary values in equation (2). One would also expect an overall decay to the oscillation. The theory assumes that the state of the GPSDC is optimally estimated, and it may be relevant that DUT3 has a longer default averaging period due to the assumption of the oscillator being more stable.




Fig. 14. Oscillator behavior when the gain controls P and I were set to high values.

Interestingly, as shown in Figs. 14 and 15, the three DUTs behaved differently in the high gain control configurations. The DOCXO-equipped DUT3 showed even more extreme behavior than the others. It exhibited large oscillations even when the P and I values were 50 or 100, while the OCXO DUTs were stable for those values. None showed a decaying exponential, while DUT2 showed a growing exponential for (P,I)=(100,100) and (50,100).




Fig. 15. Detail of previous figure, showing several forms of oscillator behavior.

The contrast between low and medium-high gain controls is shown in Fig. 16. Evidently the higher gain control values are better able to initially remove the effects of the temperature variations.




Fig. 16. DUT behavior during 3 °C temperature variations with (P,I) first (50,50) and then (5,5). The DUT3 data are not shown for the times they oscillatory.

The jitter is also a function of the gains. Fig. 17 shows the effect of altering the P and I gain controls. It is evident that high P and I values lead to less jitter in the OCXO data and more jitter in the DOCXO data, as measured by the 2-hour averages of the first differences. However, the highest-allowable values (P,I)=(500,100) show poor results for both types.




Fig. 17. Frequency of DUT data as found by averaging first differences.

IV. CONCLUSIONS AND RECOMMENDATIONS

It is evident that GPSDCs, even if from the same manufacturer, are not identical. It is also shown that the analytic treatment for PID controllers may not always apply, perhaps due to the details of the phase-locking techniques, which are often proprietary. We therefore suggest that the user who is interested in the short-term response to temperature fluctuations or other disturbances test the units. Since thermal chambers can be expensive, the user may consider rapidly switching a delay line into the antenna feed, which would mimic a temperature change in the oscillator regarding how the control loop steers it back to a setpoint.

It was not surprising to find that the DOCXO was almost always far less sensitive to temperature fluctuations, although in instances of high gain controls it displayed oscillatory behavior when the other units were stable. The factory settings for the P and I gain control values are lower for the DOCXO than the OCXO, so the manufacturer is taking this into consideration already, although our findings show that larger than default values improve the performance during temperature changes for both types. For the OCXOs, the assumption of linearity in a GPSDC's temperature dependence is largely but not entirely valid. For best performance the compensation factor should therefore either be determined in the regime wherein the unit will be operated or modelled with a non-linear function.

V. DISCLAIMER

As a matter of policy, NIST does not endorse commercial products and identifies them only if necessary for technical clarity. In this case we also caution that, since the purpose of this experiment was to evaluate steering options of products, some of their important features were disabled. It's also very unlikely that a GPSDC will be subjected to 40 °C temperature variations in short periods.

REFERENCES

- [1] D. Matsakis, "Controlling Clocks with PID Controllers", Proceedings of ION-PTTI, January, 2020.
- [2] F.M. Gardner, "Phaselock Techniques," Third Edition, Wiley Press, 2005.
- [3] V.F. Kroupa, "Phase Lock Loops and Frequency Synthesis," John Wiley & Sons, Ltd., 2003.
- [4] J.W. Clark, D.N. Matsakis, A.N. Novick, and M.A. Lombardi, "Establishing Traceability with a GPS Disciplined Clock: A Preliminary Report", International Telemetry Conference, Phoenix, Az, 2022.
- [5] A. Novick, M. Lombardi, D. Matsakis, and J. Clark, "Measurement of Transient Environmental Effects in GPS-Disciplined Oscillators", Proc. ION-PTTI, Long Beach, Ca. 2023.