App_51

Synchronizing a

One Pulse Per Second (1PPS) output

to UTC with High Precision

Precise Time and Frequency, LLC, 50L Audubon Road, Wakefield, MA 01880, USA
Tel: (+1) 781 245 9090 Fax: (+1) 781 245 9099 www.ptf-llc.com

Precise Time and Frequency, LLC

Table of Contents

1. Executive Summary 3
2. Introduction 4
3. Analysis 4
4. Improving the 1PPS/UTC Synchronization 5
5. Implementation 7
6. Technical Approach 8
7. Control Loop Calculations 8
8. Validation of Calculations 12
9. Performance Evaluation 14
11. Test Results 16
12. Conclusion 24

Precise Time and Frequency, LLC

1. Executive Summary

A special development program was undertaken at Precise Time and Frequency, LLC (ptf) to determine the feasibility of using a GPS receiver to provide a one pulse per second (1PPS) output synchronized to the international time standard, Universal Time Coordinated (UTC) within an accuracy of 10 nano seconds or less.

The development consisted of specially designed GNS receiver hardware based on the ptf 3207A with an internal rubidium atomic oscillator and temperature control, together with purpose written software utilizing a specially devolped control algorithm optimizing 1PPS synchronization over frequency accuracy/stability.

The resultant hardware and software was tested over a period of several months, and yielded a demonstrable one sigma synchronization accuracy of <6.5 nano seconds over a one day period, and better than 10 nano seconds over a 14 day period with measurement offset from the caesium standard removed.

The main limitation for demonstrating performance turned out to be the long term stability of the high performance caesium used as a reference. The caesium specification for frequency accuracy is <5E-13, which equates to a time error over 14 days of approximately 600 nano seconds. In practice it appeared that the frequency offset of the caesium standard was around 7E-14, which over a 14 day period would potentially have contributed approximately 84 nano seconds to the synchronization error.

In conclusion it was determined that the instrument performed at least as well as expected with the probablility that with a more accurate reference, better results could be demonstrated.

PTF
Precise Time and Frequency, LLC

2. Introduction

In todays world of high speed data transmission precision frequency references abound. High performance references are required to drive the various system elements transmitting/receiving the data such as satellite modems, upconverters, and downconverters.

Typically the frequency references are designed for extremely high accuracy (<1E-12), excellent stability (<1E-11 at 1 second), and very low phase noise/jitter (<-130dBc at 10Hz offset) to insure maximum data throughput for the available bandwidth. The result is that everything tends to be optimized for the ultimate frequency performance, however when considering the 1PPS synchronicity to UTC this may not be ideal.

In this paper, a technique is described whereby the 1PPS synchronization to UTC is maintained to very small variations, with minimal sacrifice of the frequency accuracy and stability.

3. Analysis

Consider the following:

A frequency reference that is stable and accuracte to 2E-12 has a fixed time offset drift as follows:

From the basic equation relating the change in time offset to frequency offset:

$\frac{df}{F} = \frac{dt}{T}$ where df is the frequency error(Hz), F is the carrier frequency(Hz), dt is the time error(seconds), and T is the time over which the time error is measured(seconds). df / F is usually referred to as the fractional frequency error.

In this case: $\frac{df}{F} = 2E\text{-}12$ this means that for a nominal desired frequency of 10,000,000 Hz we may have an error resulting in an actual frequency of 10,000,000.00002 Hz

If we now consider a time period for measuring the error of say 1 hour then:

$$2E\text{-}12 = \frac{dt}{3600} \quad \text{(dt and T are in seconds)}$$

Therefore $dt = 2E\text{-}12 \times 3600 = 7.2\ E\text{-}9$ i.e. 7.2 nano seconds

In words this means that a frequency standard operating at an accuracy of 2E-12 may drift in phase by approximately 8 nano seconds in 1 hour. This error would be passed to the 1 pulse per second (1PPS) output which is directly derived from the frequency reference, usually 10MHz (divided by 10,000,000 to give the 1PPS output), so that the 1PPS error would be increasing at a rate of 8 nano seconds per hour, theoretically reaching an offset of 96 nano seconds after 12 hours.

PTF
Precise Time and Frequency, LLC

In practice, for most frequency references, this does not happen, as the methodology for controlling the frequency by default is designed to minimize the 1PPS offset, albeit over long periods of time. The preferred method is typically to use a 1PPS reference from satellites which are synchronized to UTC (GPS for example). This satellite derived 1PPS is averaged over a fairly long period of time (100 seconds or more) to reduce the short-term noise, and then compared to a 1PPS derived from an internal (typically 10MHz) frequency reference.

The phase error between the two 1PPS signals (satellite derived and internally derived) is driven to zero by use of a phase lock loop (PLL), and this in turn maintains the frequency accuracy. The compromise, between synchronization to UTC and a highly stable and accurate frequency reference, comes because as the satellite signal contains a lot of short-term noise, the satellite derived 1PPS needs to be averaged over as long a period as possible for best frequency stability and accuracy.

The result is that time constants for controlling the internal frequency generator (e.g. high performance OCXO or rubidium atomic oscillator) are typically very long (100 to several thousand seconds), therefore driving the frequency reference to give a zero error could take several hours. This can result in a peak short term 1PPS/UTC synchronization error of tens of nano seconds.

For this reason, most high-performance frequency references quote a performance specification for 1PPS/UTC synchronization in terms of a statistical measurement as 30 nano seconds, 1 sigma. This means that statistically speaking, if you were to measure 100 samples of the 1PPS synchronization to UTC, then 68 (68%) of the readings should be within 30 nano seconds.

4. Improving the 1PPS/UTC Synchronization

In order to improve the synchronization accuracy of 1PPS to UTC it is necessary to reduce the system control loop time constant such that the error of the internal PPS relative to averaged satellite PPS becomes the focus.

Determining the ideal 1 PPS synchronization control loop time constants is approached as follows:

Determine the maximum 1PPS/UTC synchronization error required. For this example, we will use 10 nano seconds as the desired maximum time error, and make use of the standard equation that describes the relationship between frequency offset and time error:

$$\frac{df}{F} = \frac{dt}{T}$$

For a good frequency reference standard, the fractional frequency error would typically be around 2E-12. Unlike the “good old analog days” of continuous variability, the control loop used to “lock” the PLL will invariably have a finite resolution caused by the conversion from digitally calculated values to analog

implementation. In the case of an OCXO this will be the minimum step size of the D/A converter, in the case of a rubidium this will probably be the minimum step of a DDS. In the latter case this would typically be around 7E-13. This will define the maximum loop time constant of the control loop compatible with the desired 10 nano second maximum PPS error for synchronization

Verify the elapsed time to reach our maximum desired PPS error of 10E-9 given the frequency control resolution of 7E-13;

$$df = \frac{dt}{F} = \frac{7E\text{-}13}{T}$$

$$\text{Therefore } T = \frac{dt}{7E\text{-}13}$$

In this case, our maximum 1PPS time error, dt, is ten nano seconds (10E-9), therefore, assuming that dt starts at 0, the elapsed time (T) before dt reaches 10E-9 (10 nano seconds) with a frequency offset (error) of 7E-13 is:

$$T = \frac{10E\text{-}9}{7E\text{-}13} = \frac{10E4}{7} = 14{,}286 \text{ seconds (approximately 4 hours)}$$

This figure of course assumes that everything else within the system is absolutely stable, not really a practical expectation!

Some of the parameters can be better controlled than others e.g. controlling temperature will decrease variations caused by temperature change, probably the biggest culprit.

Other variations are less controllable, such as variations from satellite to satellite as satellites recede from view and new satellites come into view.

PTF
Precise Time and Frequency, LLC

5. Implementation

From the above, the theoretical limitation on maintaining synchronization to within 10 nano seconds with an initial frequency accuracy of 7E-13 is 14,000 seconds. As far as the other perturbations go, temperature is probably the most significant, and fortunately also the most controllable.

The implementation described here utilized a PID temperature controller, controller an instrument fan. Also the entire instrument was housed within a 19 inch rack, at a room temperature of a nominal 22 degrees Celsius. The inside temperature rise of the equipment was approximately 35 degrees, and therefore the set point of the controller was set to 45 degrees, +/- 0.2 degrees.

To give some margin, the time to maintain synchronization to 10E-9 with an initial offset of 0, and an accuracy of 7E-13 was assumed to be 10,000 seconds.

Using the rule-of-thumb guideline of five time constants (99% of error is removed after five loop time constants), this would give a loop time constant of 2,000 seconds, i.e. a loop bandwidth of around 0.5 milli Hz.

Also, to ensure that drift of the frequency source was minimal over the time period, a rubidium internal oscillator was used, with an (uncontrolled) aging rate of 5E-11 per month.

To determine what this aging rate represents per second:

5E-11 / month = 5E-11 / 30(days) x 24(hrs) x 3,600 (secs) = 5E-11 / 2,592,000 secs

5E-11 / 2,592,000 = 1.929E-17 frequency change in one second

In time terms over 2000 seconds this represents:

from $\frac{df}{F} = \frac{dt}{T}$ 1.929E-17(df/F) x 2000 (T in seconds) = 3.86E-14 (t in seconds)

or 3.86 E-5 nano seconds, which relative to our desired <10 nano seconds can be considered negligible.

PTF
Precise Time and Frequency, LLC

6. Technical Approach

With reference to the above calculations, for best stability it would appear that we should minimize any perturbations from the control loop for the maximum time period possible, compatible with our 7,000 to 10,000 estimated limit for maintaining the 1PPS within <10 nano seconds to UTC.

With this in mind therefore, an internal rubidium atomic oscillator was used and a two term (proportional and integral) control loop developed, detailed calculations are shown below.

7. Control Loop Calculations

Characterization of the internal rubidium oscillator

The rubidium has an "integration" within the unit where the frequency control has the following time domain difference equation:

$$y_n = x_n + y_{n-1} \quad (1)$$

The above equation is not a pure integrator.

$$Y(z) = X(z) + Y(z)z^{-1}$$

$$\frac{Y(z)}{X(z)} = \frac{1}{1 - z^{-1}}$$

Applying the bilinear transform where T is the sampling interval (T = 1sec in this case), we can show that the result is not simply 1/s:

$$z = \frac{2 + Ts}{2 - Ts}$$

$$\frac{Y(s)}{X(s)} = \frac{2}{Ts} + 1 = \frac{2}{s} + 1$$

To remove the "integration" multiply the overall transfer function by:

$$1 - z^{-1}$$

The difference equation from above is:

$$a_n = x_n - x_{n-1}$$

Precise Time and Frequency, LLC

Substituting a_n into x_n in the rubidium equation above (1) to see if the cancellation occurs we get:

$$y_n = x_n - x_{n-1} + y_{n-1} \quad (2)$$

$$Y(z) = X(z) - X(z)z^{-1} + Y(z)z^{-1}$$

$$Y(z)(1 - z^{-1}) = X(z)(1 - z^{-1})$$

$$Y(z) = X(z)$$

Cancellation does occur and y_n equals x_n although it is not intuitively seen in the difference equation (2).

Lag-lead compensation filter.
The lag-lead compensation filter is as follows:

$$y_n = y_{n-1} + k_1 x_n - k_2 x_{n-1}$$

Substituting a_n into x_n of the previous equation:

$$a_n = x_n - x_{n-1}$$

$$y_n = y_{n-1} + k_1(x_n - x_{n-1}) - k_2(x_{n-1} - x_{n-2})$$

$$y_n = y_{n-1} + k_1 x_n - k_1 x_{n-1} - k_2 x_{n-1} + k_2 x_{n-2}$$

$$y_n = y_{n-1} + k_1 x_n - (k_1 + k_2)x_{n-1} + k_2 x_{n-2} \quad (3)$$

Equation (3) is the lag-lead filter with the cancelling effect of the rubidium "integration". Coefficients k_1 and k_2 are either both positive or both negative depending on the control sense.

Multiplying the transfer function of (3) by $(1 - z^{-1})$ should result in the original lag-lead equation:

$$Y(z) = Y(z)z^{-1} + k_1 X(z) - (k_1 + k_2)X(z)z^{-1} + k_2 X(z)z^{-2}$$

Simplify

$$\frac{Y_o(z)}{X_o(z)} = \frac{k_1 - (k_1 + k_2)z^{-1} + k_2 z^{-2}}{(1 - z^{-1})}(1 - z^{-1})$$

Precise Time and Frequency, LLC

$$\frac{Y_o(z)}{X_o(z)} = k_1 - (k_1 + k_2)z^{-1} + k_2 z^{-2} = \frac{k_1 - k_2 z^{-1}}{1 - z^{-1}}$$

The validity of the above equation was tested by writing and running some code for the model:

```
// Phase Lock Rubidium Loop Using 3207A
// K. Lyon 12/3/2019   Loop BW = 1mHz.

#include "stdio.h"
#include "math.h"

main(void) {

        double a1 = 0.969;          // Pre-filter coefficients
        double a2 = 0.015;

        double k1 = 11.727;        // Lag-Lead filter coefficients
        double k2 = 11.675;

        double transient;
        double G = 7.6e-4;         // Overall DC gainterm

        double y, prev_y, x, prev_x;
        double y_filter[2000];
        double prev_y_comp, y_comp;
        double y_int, prev_y_int;
        double a, prev_a,b;
        int i;

        int count = 2000;

        // Initial Conditions
        x = 0 = y_filter[0] = y_filter[1] = prev_x = x = 0 ;

        prev_y = y = prev_y_comp = 0;
        transient = 0;
        y_int = prev_y_int = prev_a = 0;

        FILE *ptr;
```

Precise Time and Frequency, LLC

```
        printf("  Time      x\n");

        ptr = fopen("PLL.dat", "w");
        fprintf(ptr,"Time     x\n");

        for(i=2; i<=count; i++) {
                printf(" %d     %e\n",i,x);
                fprintf(ptr,"%d    %e\n",i,x);

                if(i >10)
                transient = 100;

                x = G*(x + transient);

                // Low pass filter
                y_filter[i] = a1*y_filter[i-1] + a2*(x + prev_x);
                prev_x = x;

                // Lag-Lead Compensation + Intergation Cancellation
                        y_comp = prev_y_comp + k1*y_filter[i] - (k1+k2)*y_filter[i-1] + k2*y_filter[i-2];

                // Pure Integration (frequency to phase) from phase detector
                y_int = prev_y_int + 0.5*(y_comp + prev_y_comp);

                prev_y_comp = y_comp;
                prev_y_int = y_int;

                        // "Integration" of Rb control word by Rb
                        a= prev_a + y_int;
                        prev_a = a;

                        x = -a;  // close the loop
        }

        fclose(ptr);
        return(0);
}
```

Precise Time and Frequency, LLC

8. Validation of Calculations

Transient response result for Loop BW = 1mHz.

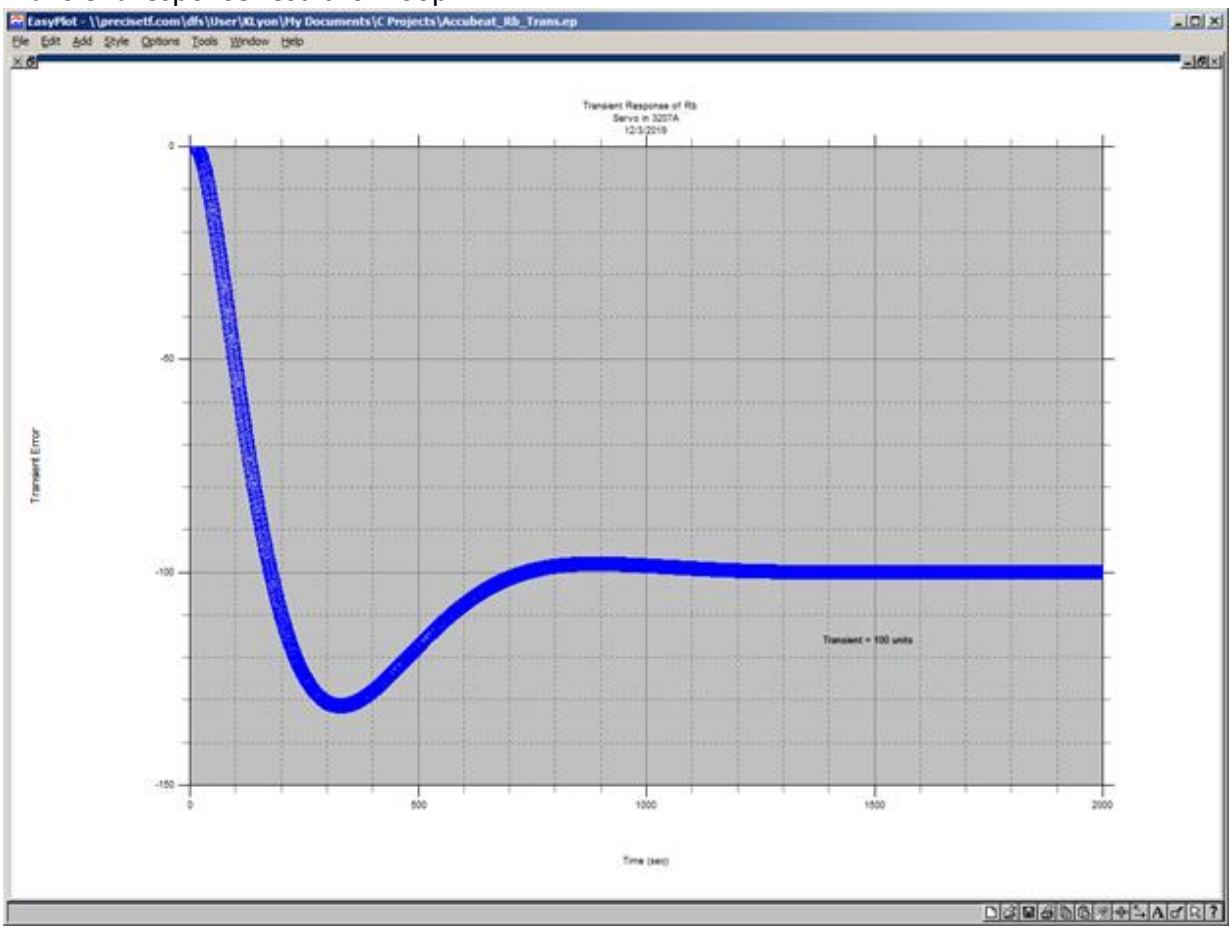

Calculation Summary

For loop bandwidth of 1mHz:

$$y_n = y_{n-1} + k_1 x_n - (k_1 + k_2) x_{n-1} + k_2 x_{n-2}$$

For low-pass filter: $y_n = a_1 y_{n-1} + a_2 (x_n + x_{n-1})$

For: $k_1 = 11.727$ $k_2 = 11.675$ $a_1 = 0.969$ $a_2 = 0.015$ $\varsigma = 0.707$

The above equation cancels out the rubidium "integration" of the control word.
$\sigma(\tau)$ peaks at about 200sec then drops at a WPM rate.

To reduce the peaking seen above a loop with more damping was tried:

$k_1 = 33.12765$ $k_2 = 33.0756$ $a_1 = 0.96904$ $a_2 = 0.01548$ $\varsigma = 2.0$

The above coefficients moved the peak to about 100sec, indicating that the MathCad analysis was not quite correct.

It appeared that the pre-filter was affecting the transient response – therefore it was decided to push pre-filter pole to 20x the loop bandwidth (c = 20, was 5), try a1 = 0.88166 and a2 = 0.5917 with $\varsigma = 2.0$.

This still gave overshoot – but should be underdamped so it appeared that the Mathcad approximations were not working exactly, however the C computer model and MicroCap agreed very well.

To get a damped system it was decided to try getting a damped system by adjusting τ_p and τ_z in MicroCap. A very good response was obtained with $\tau_f = 2.0$, $\tau_p = 500$ and $\tau_z = 5000$. Bandwidth is approx. 1 mHz. Calculate difference equation coefficients from τ_p , τ_z and τ_f.

$k_1 = 10.0010$ $k_1 - k_2 = 0.002$,

These coefficients produce a long tail on the transient response.

$k_2 = 9.9990$ $a_1 = 0.600$ $a_2 = 0.200$

The MathCad file was updated so only τ_p. had to be provided. τ_z was calculated to provide a critically damped system (equal real roots, $\varsigma = 1$). The pre-filter pole was placed 80x from the loop bandwidth. The pre-filter pole could be put closer to loop bandwidth. The coefficients were very close to the original calculations although there was a larger difference between k_1 and k_2. If the difference becomes very small the loop may not work properly due to round off.

$\tau_p = 50$, $\tau_z = 513$, $A = 7.6(10^{-4})$ and $\tau_f = 2.0$ for the following coefficients:

$$\tau_z = 2\sqrt{\frac{\tau_p}{A}}$$

$k_1 = 10.26978$ $k_1 - k_2 = 0.020$
$k_2 = 10.24978$

$a_1 = 0.59799$ $a_2 = 0.20101$ $\varsigma = 1.0$

The transient with the above coefficients gave overshoot although according to the literature* it should not overshoot. The roots were indeed the same and the time domain solution for a step input consisted

Precise Time and Frequency, LLC

of only exponentials (it is indeed possible to get overshoot with exponentials). The closed loop gain was free of peaking around ω_n

It turned out that we were not dealing with a simple second-order equation because the numerator contained a derivative, so all of the simple second order calculations and plots in the text did not apply, however in evaluating the calculated result implemented in an instrument, performance was good.

The PPS from the GPS engine has an ADEV of about 10^{-8} at 1sec with a WPM slope. The loop causes the rubidium WFM to increase near the loop bandwidth to match the PPS WPM noise slope.

9. Performance Evaluation

The above control algorithms were implemented in a 2U version of the ptf 3207A GNS receiver with an internal rubidium oscillator and multi GNS capability.

To minimize any affects from temperature variations, a substantial heatsink was attached to the rubidium oscillator and a PID controller controlling a small fan was fitted, with a measurement thermistor attached to the rubidium heatsink. Temperature control was set to +/- 0.1 degrees at a temperature of 49C

For measurement and synchronization of the 1PPS output relative to UTC, it was necessary to have an independent and very stable alternative 1PPS source as a reference.

The alternative source was provided by a high performance cesium standard. The accuracy specification for the high performance cesium is <5E-13. As discussed above, an error of 5E-13 would result in a 1PPS phase drift of 10 nano seconds over a period of 20,000 seconds.

The 1PPS output of the instrument is derived from the internal 10MHz generated from the rubidium oscillator, and therefore to measure the stability and synchronization of the 1PPS, the 10 MHz output was used, and compared to the 10MHz output of the high performance cesium by means of a picopak time interval measurement module.

The measurement uncertainty introduced due to the inherent drift of the cesium meant that results measured over time periods of longer than 20,000 seconds (approximately 5.5 hours) would likely include a significant contribution to the error from the reference itself.

Results obtained were plotted graphically, showing the difference between reference and the unit under test, and also statistically to give an estimate of the 1 sigma standard deviation of the synchronization error.

From the results it was found that, even over longer time periods, the one sigma synchronization was well within a 10 nano second expectation.

10. Test Setup

The test setup used is shown in the illustration below and comprised a Agilent 5071ª High Performance caesium standard, a PicoPak measurement system, and the unit under test, a ptf 3207ª with internal rubidium, temperature control, and control algorithm optimized for 1PPS sunchronization.

1PPS UTC Synchronization System Test Setup

Precise Time and Frequency, LLC

11. Test Results

The following graphs and histograms were produced from data taken over a period of two weeks (1,218,105 seconds). Data was taken at one second intervals, however due to limitations of the statistical analysis software (Stable32) the data could only be plotted in sections of 1 million samples (approximately 11.5 days) at one time.

Interesting to note is that over this time period, total phase error attributable to the high performance cesium reference standard could be up to approximately 500 nano seconds.

Also, as the initial offset is arbitrary (the cesium 1PPS position is not synched to UTC) in most cases it is convenient to first normalize the phase data, to eliminate the initial phase offset. This was not done for the overall data set in the first couple of plots

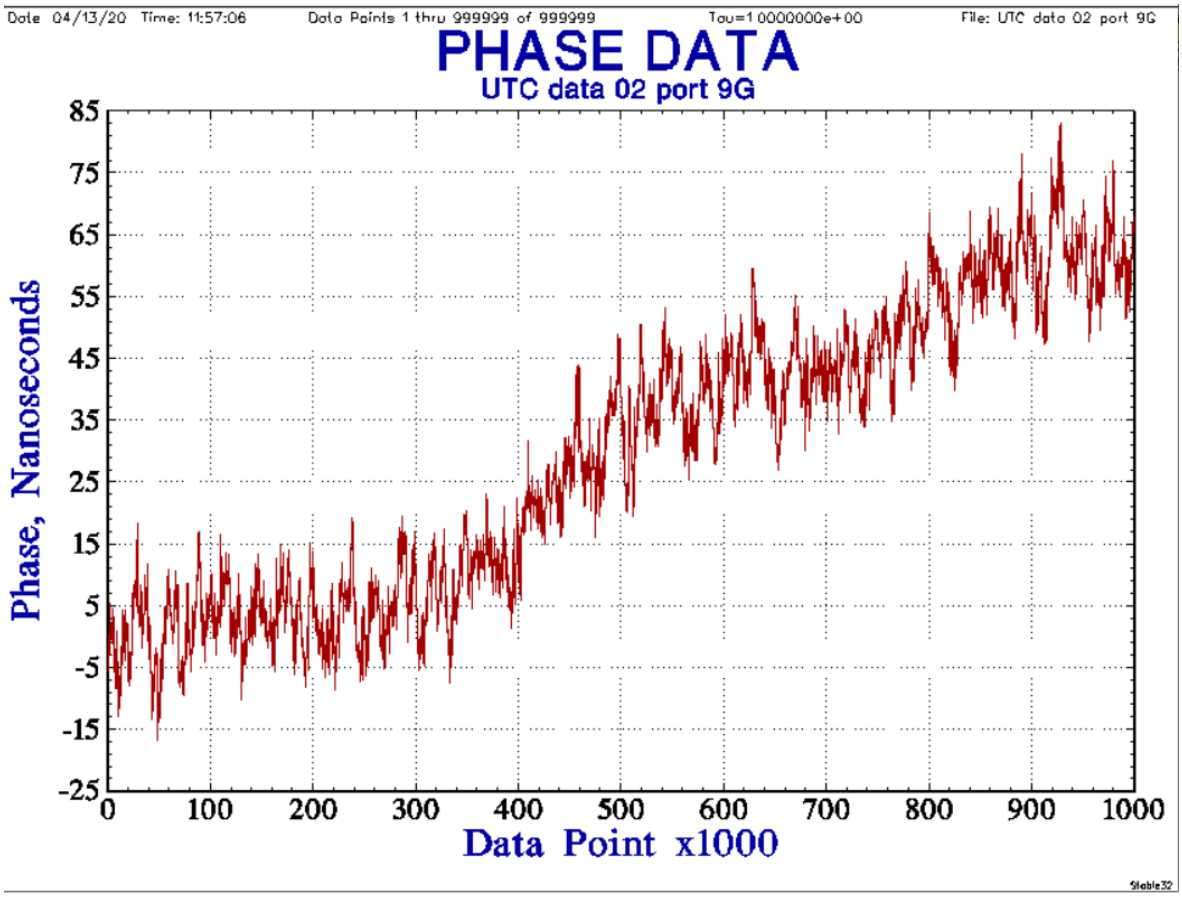

Phase plot of the first 1E6 Samples (Not normalized)

The above plot shows that whilst the initial period of approximately 300,000 seconds demonstrates a total variation of around 10 nano seconds, as the data continues there is a phase "drift" that is probably due to the cesium reference of approximately 60 nano seconds over a period of approximately 700,000 seconds, which equates to a frequency error of approximately 5E-14 on the cesium, 10 times better than the specification of 5E-13

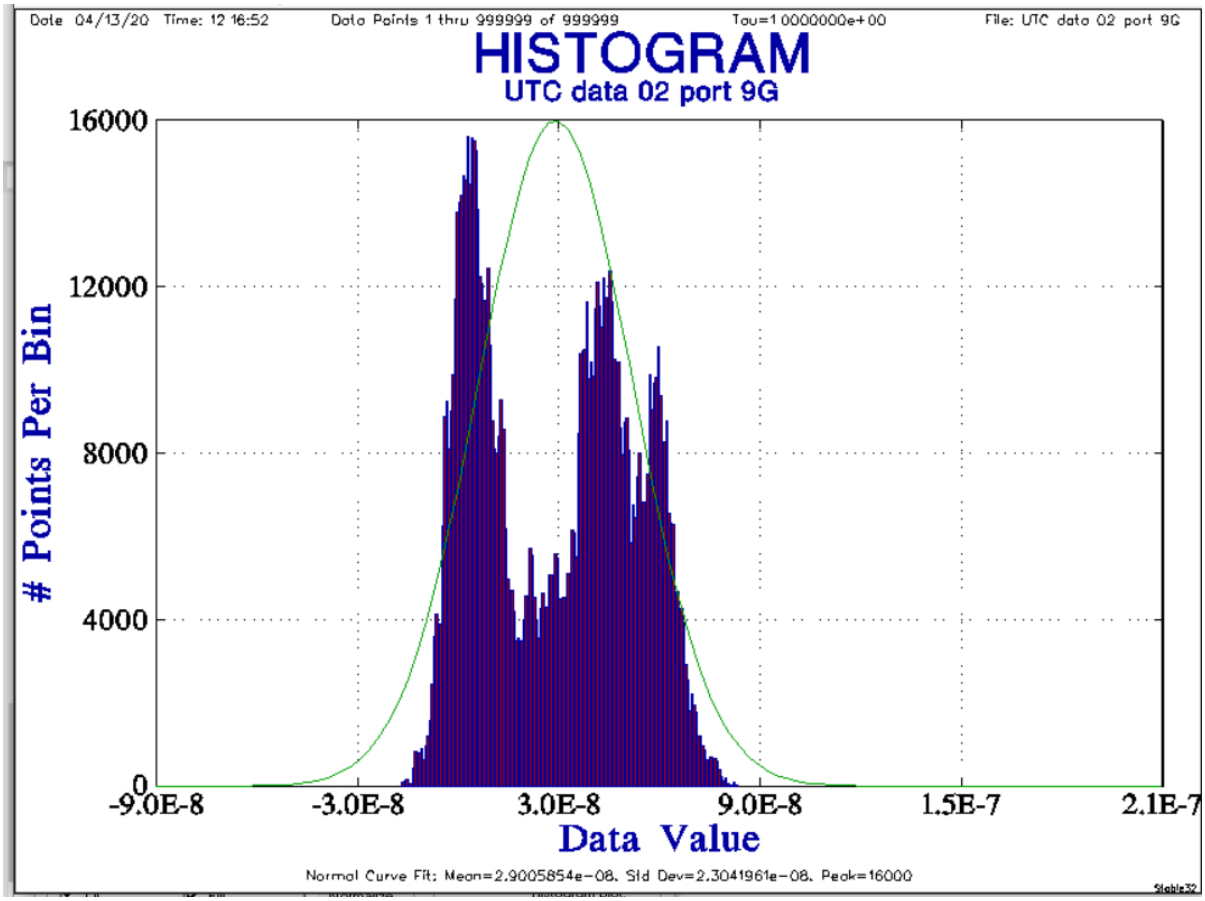

Histogram Plot of the first 1E6 samples

The above histogram plot shows that over 1E6 seconds (11.5 days) total variation was approximately 40 nano seconds, with a statistically calculated 1 sigma value of 23.04 nano seconds.

The next sets of data are more representative of actual performance and are plotted from data taken over periods of 100,000 (1E5) seconds, which with the measured performance of the cesium standard would give an error due to our reference of approximately 5 nano seconds, calculated from:

$$\frac{df}{F} = \frac{dt}{T} \qquad \text{or} \quad df \times T = dt \qquad \text{so: } dt = 5E\text{-}14 \times 1E5 = 5E\text{-}9 \text{ (i.e. 5 nano seconds)}$$

PTF
Precise Time and Frequency, LLC

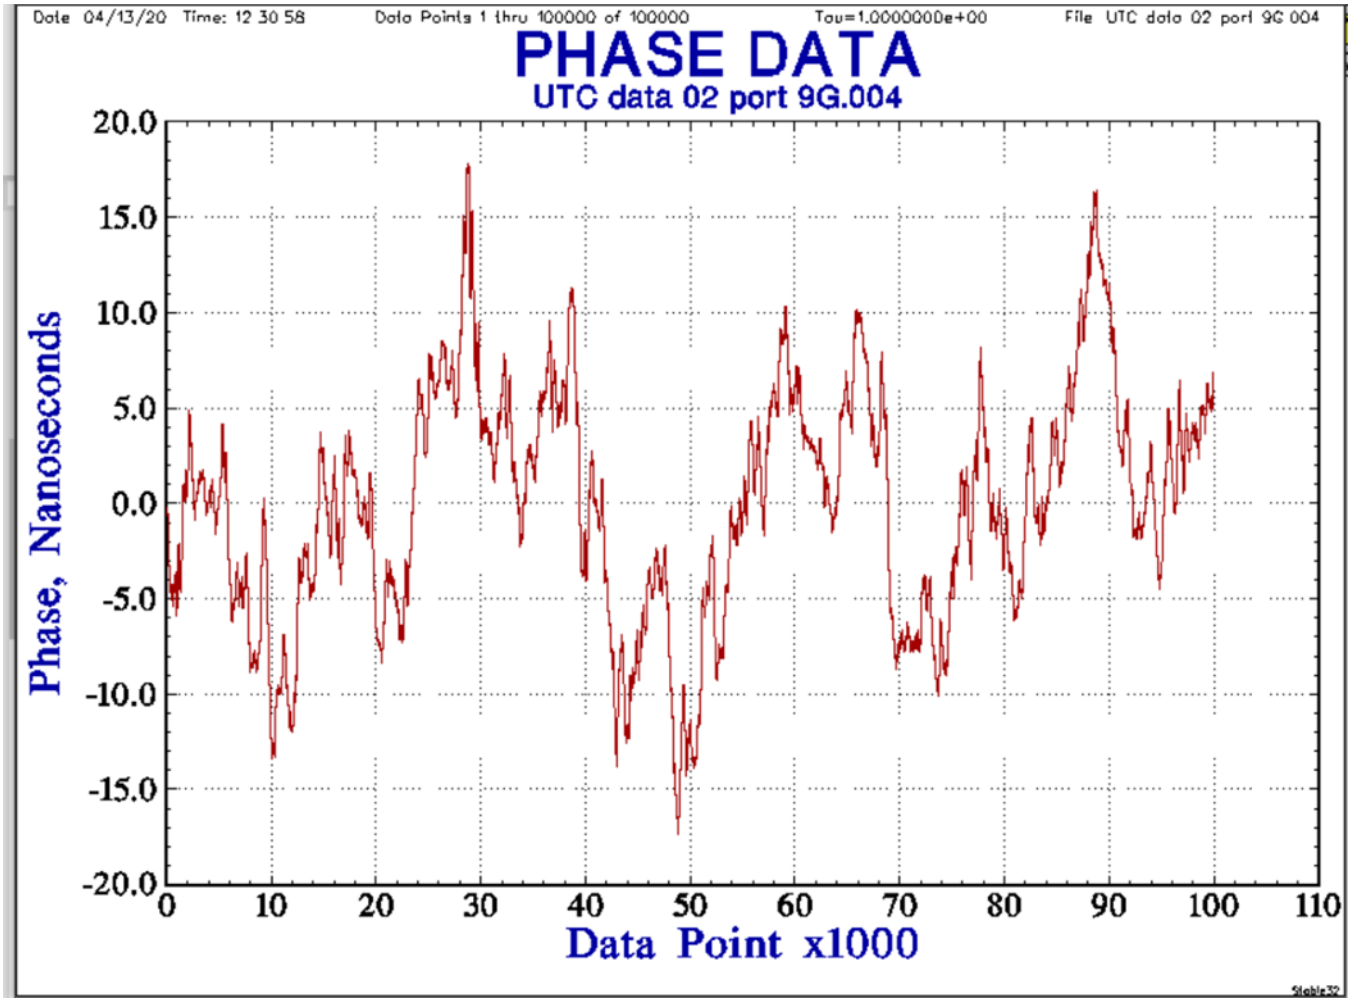

Phase plot of the first 1E5 samples – normalized to eliminate random offset

PTF
Precise Time and Frequency, LLC

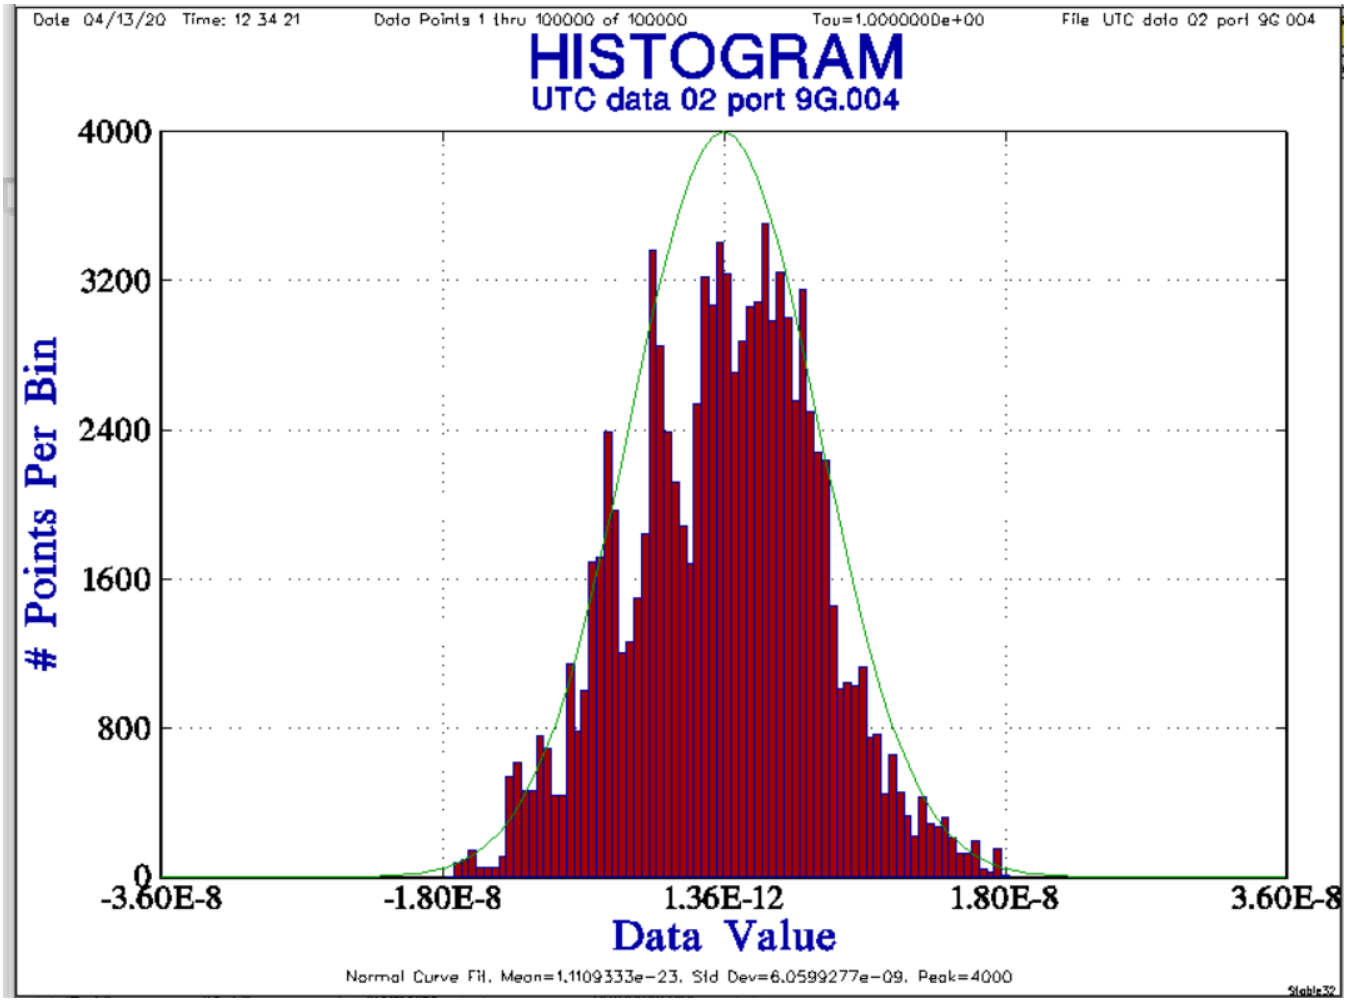

Histogram plot of the first 1E5 samples – normalized to eliminate random offset

Over a period of 1E5 (100,000) seconds, or 1.15 days, the data shows a 1 sigma accuracy of 6.05 nano seconds. When considered against the inherent drift of the reference this definitely represents a variation in the region of our intended accuracy of <10 nano seconds.

Precise Time and Frequency, LLC

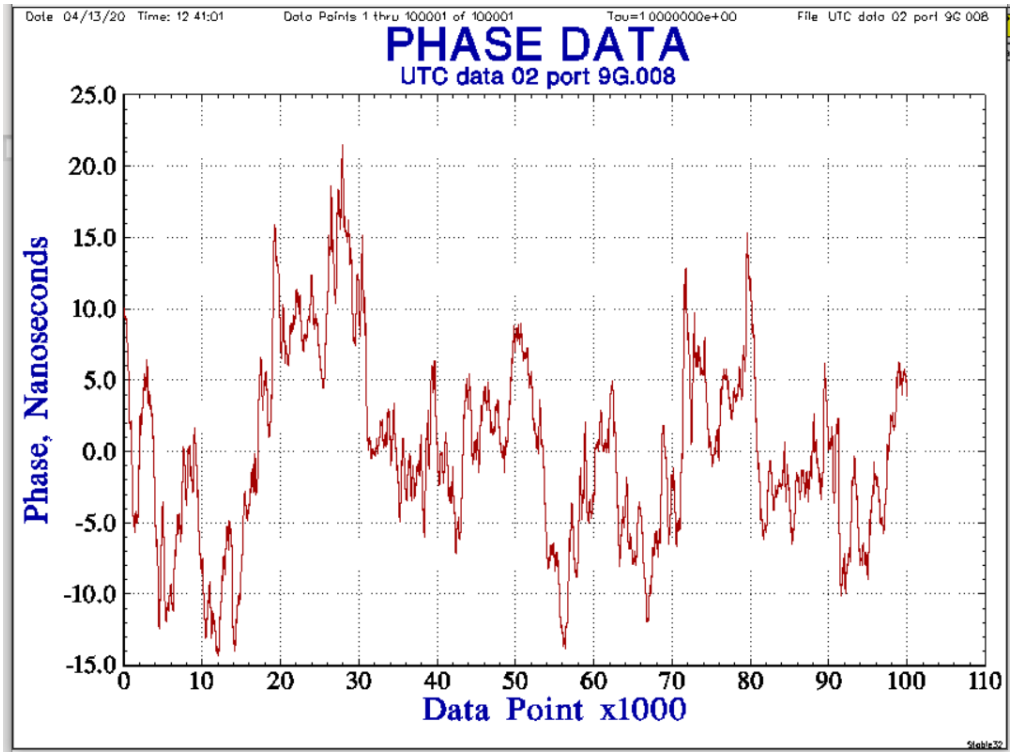

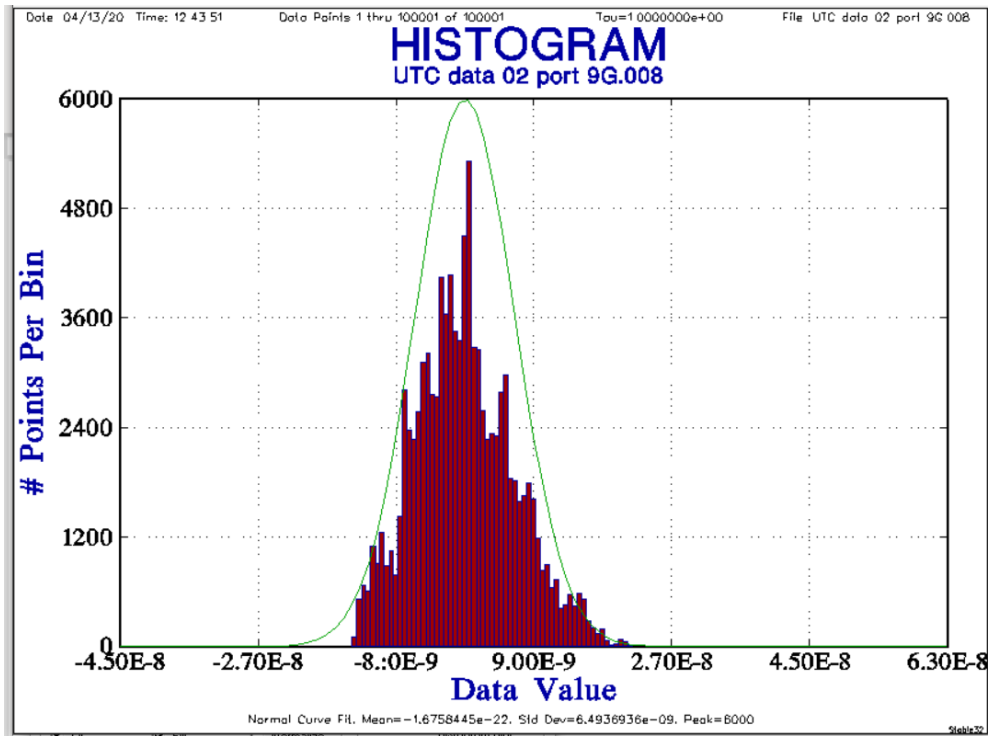

Phase plot and Histogram for samples 900,000 to 1,000,000

The above plots show that the performance once normalized is very similar to the results for the first 1E5 samples with a standard deviation of around 6.49 nano seconds.

Precise Time and Frequency, LLC

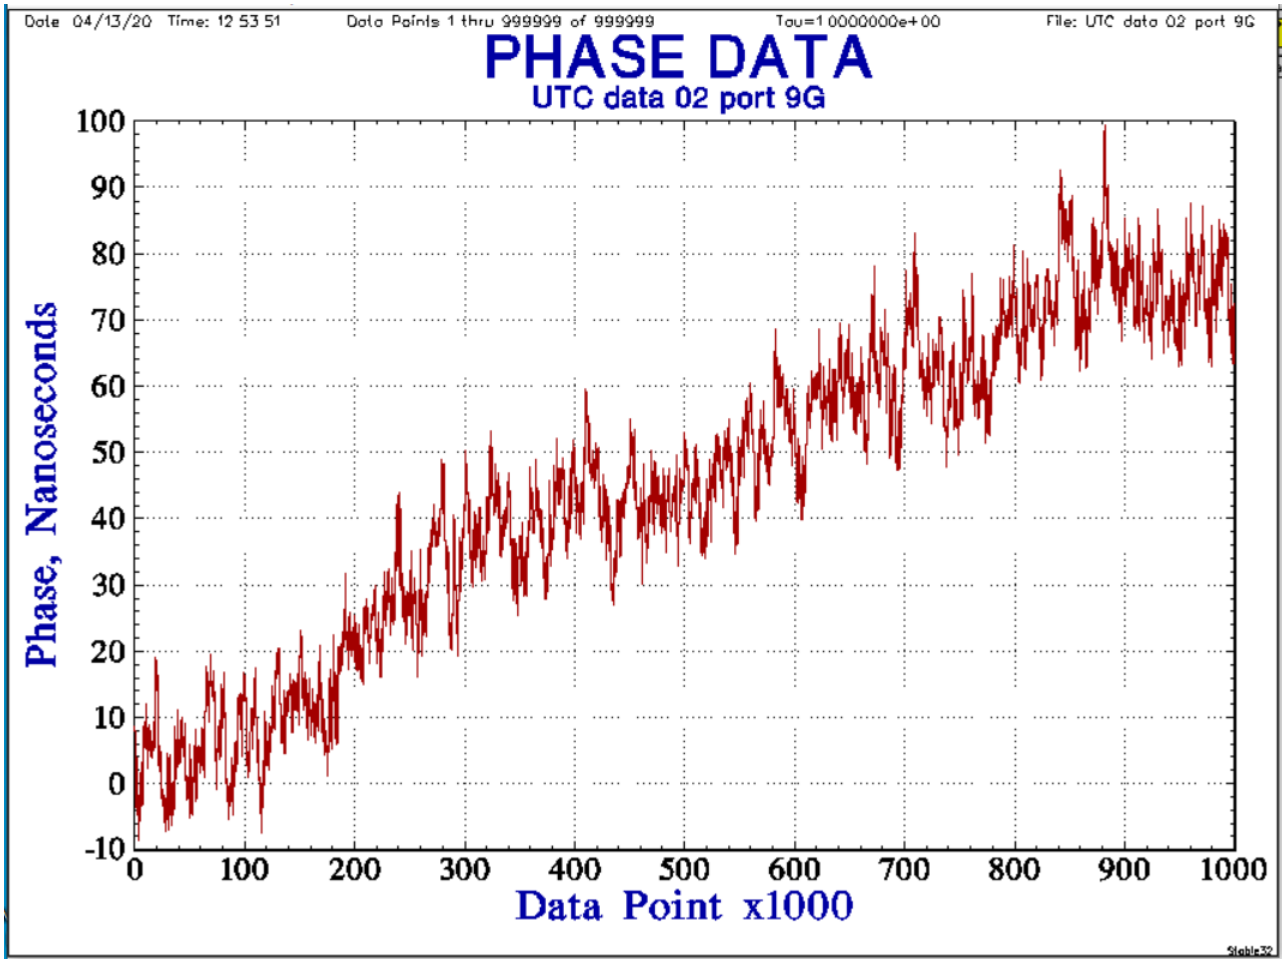

Phase plot of the last 1E5 samples – before normalization

The last set of 1E6 samples shows a continuing frequency offset trend from the first set of 1E6 samples, with a drift of approximately 70 nano seconds (7E-8)over a period of 1E6 seconds. Using our formula this gives our frequency offset at 7E-14:

$$\frac{df}{F} = \frac{dt}{T} \quad \text{or} \quad \frac{dt}{T} = \frac{df}{F} \quad \text{so:} \quad \frac{7E\text{-}8}{1E6} = \frac{df}{F} = 7E\text{-}14$$

This was highly likely due to the caesium standard having a frequency offset, the specification for which is 5E-13, so in practice 7E-14 is well within the caesium specification.

Plotting a histogram of this gives the graph below:

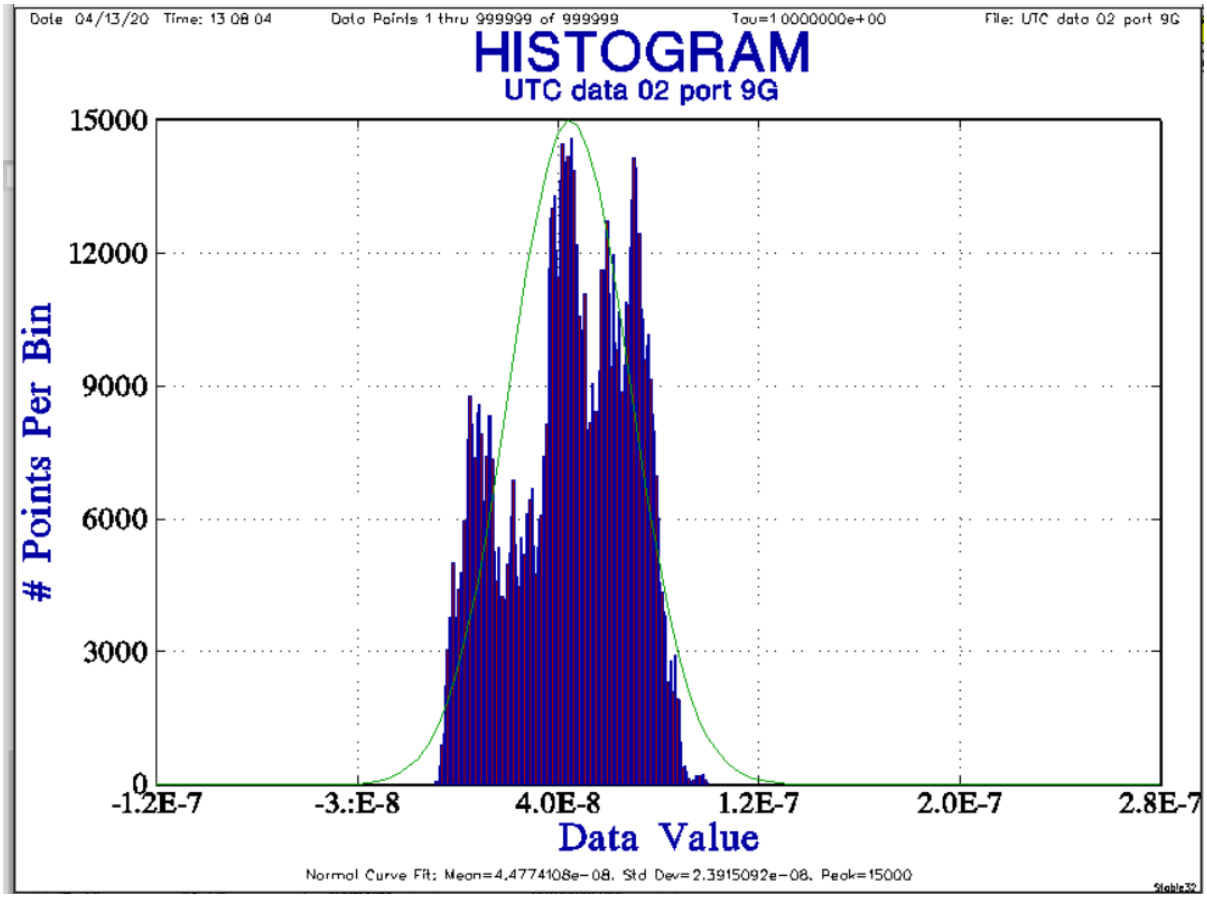

The 1 sigma standard deviation was very similar at 23.9 nano seconds (versus 23.04) again well within the error contribution that could be expected from the cesium frequency offset.

Precise Time and Frequency, LLC

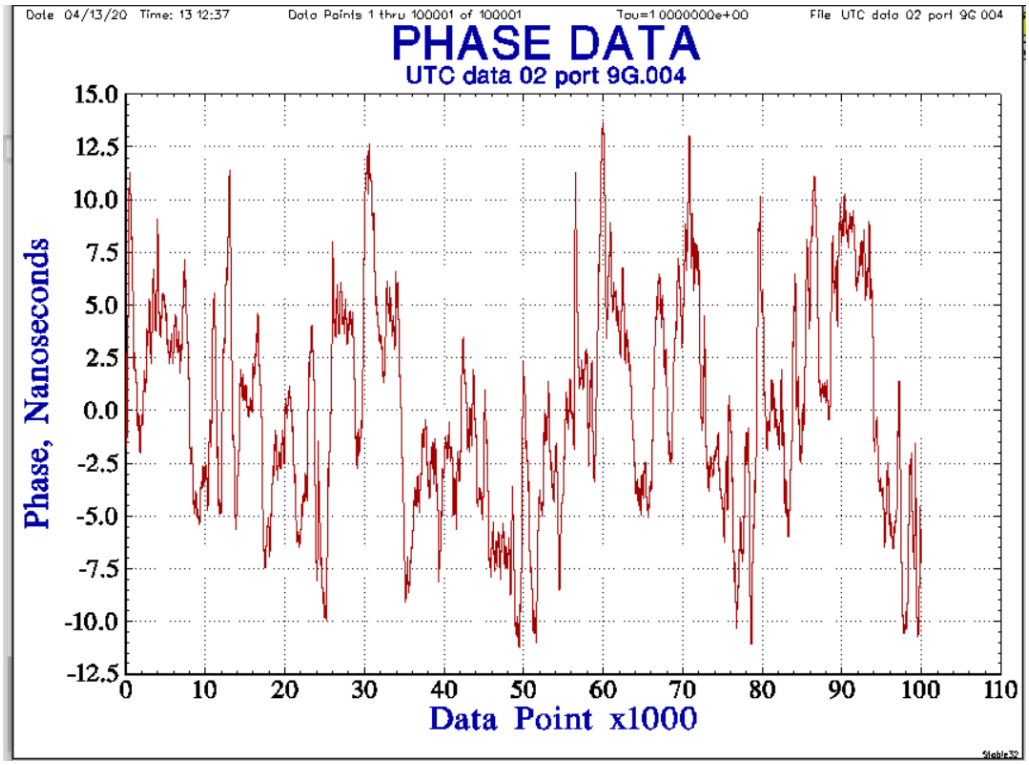

Phase plot of the last 1E5 samples – normalized to eliminate random offset

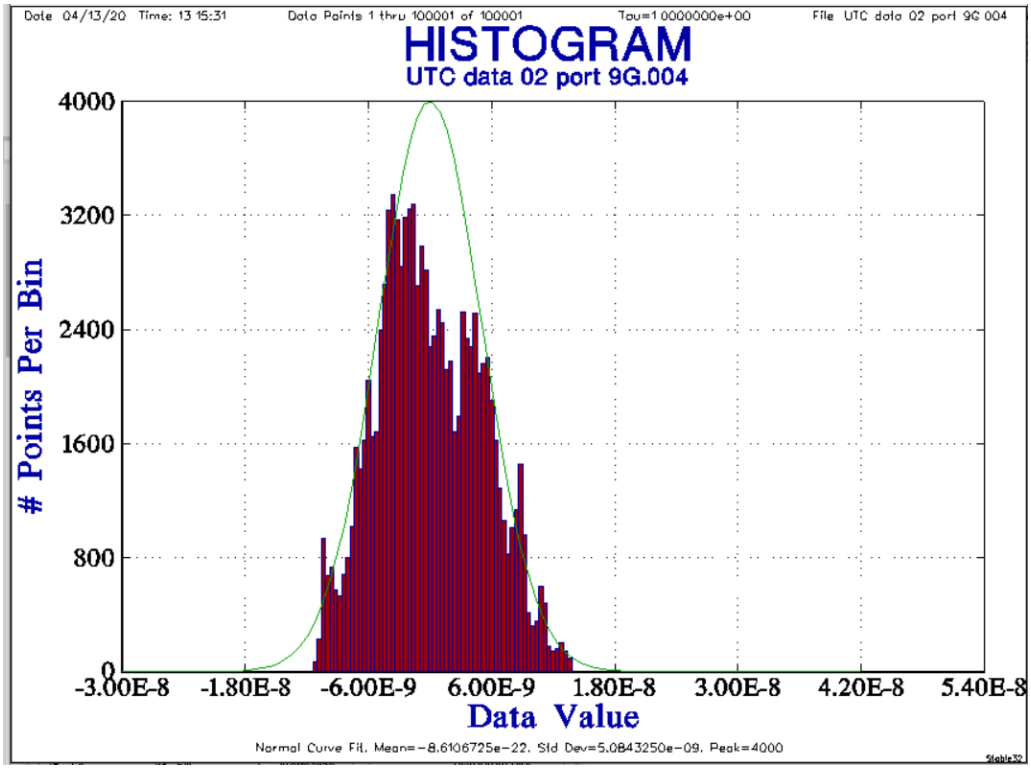

The histogram for the last 1E5 samples shows a similar result at 1 sigma of 5.08 nano seconds.

Precise Time and Frequency, LLC

12. Conclusion

With continuous monitoring and temperature control to approximately +/- 0.1 degrees Celsius, as far as could be determined within the limits of the measurement system available, the ptf 3207A with integrated rubidium was able to maintain a synchronization accuracy relative to UTC of better 10 nano seconds on a 1 sigma basis, over an extended period of two weeks.

At the same time the unit was still providing a highly accurate and stable 10MHz frequency output to within <1E-12.

For further investigation and verification, it is recommended that the system be tested against an even more stable reference such as an Active Hydrogen Maser, with the Maser in a temperature controlled environment.