Teltest Electronics Laboratories, Inc.
Austin, Texas

Calibration and Monitoring of Frequency Standards - Phase Method -

Rev 0.08

Jim Satterwhite K4HJU

A Work in Progress

Teltest Electronics
1/16/2014

Table of Contents

Table of Figures

Table of Tables

1. Introduction

In this article I am speaking primarily to the engineer, technician or experimenter who has interest in precision frequency standards and their calibration and monitoring. This article will deal with the calibration and monitoring of the various 10 MHz frequency standards.

Years ago, I acquired an HP106B frequency standard, an early '60s frequency standard intended as a reference for calibrating other standards at remote locations. It would be calibrated and put on an airplane and taken to the remote location. Then the remote standard would be calibrated using the HP106B as the reference. The HP106B contains a 5 MHz crystal oscillator in a double oven. Technology has come a long way since then.

I searched for ways to verify its calibration. What is involved and how do I go about doing it? I had designed and build some, what I consider, quite good and versatile frequency counters and was familiar with frequency measurement; at least to parts in 10^{-8} (.01 PPM - (PPM - parts per million)). When the object becomes accuracy in the vicinity of 10^{-11} or 10^{-12} (PPT - parts per trillion), the challenge increases considerably. The HP106B specification is a stability of +/- 50 PPT over 24 hours.

Figure 1 HP 106B

Low-tech ways of calibrating the HP106B would involve comparing the device under test (DUT) with a standard reference:

1. Listening on a receiver for a zero-beat signal when the output of the generator is mixed with the signal from an external reference such as WWV.
    1. Accuracy is poor, because the ear can only detect errors of a few parts per hundred.
    2. Accuracy is also determined by the low-frequency response of the audio portions of the receiver because phase errors of one or two Hertz must be audible to the observer.
2. Directly measuring the frequency of the generator with a frequency counter.
    1. Accuracy is limited because most counters have internal references which are accurate to a few parts per million (1 - 100 x 10^{-6}). See Table 1.
    2. Since the HP106B is capable of precision approaching a few tens of parts per trillion (50 PPT), the reference to which it will be compared must be at least as accurate, and by preference it will be at least an order of magnitude better than this.

2. Reference Sources

The first part of frequency calibration is to compare the unknown frequency to a frequency known to an accuracy and precision greater than the desired calibration level. That means at the outset you have to have access to a source more accurate than the desired calibration level.

References for frequency measurement cover the gambit. On the low end to stand-alone crystal oscillators to Cesium standards and beyond on the high end. Somewhat of an overview is given below for the various reference technologies.

Table 1 Relative Accuracy of Various Frequency References

Reference	Modification	Accuracy Range
Crystal	-	1 - 100 PPM
Crystal	Temperature Compensated TCXO	~ 0.1 PPM
Crystal	Ovenized OXCO	0.001 - 0.1 PPM
Crystal	Double Oven	~ 50 PPT
Crystal	GPS supervised OXCO	~ 5 PPT
Rubidium	-	~ 50 PPT
Rubidium	GPS supervised	~ 5 PPT
Cesium	-	0.01 - 0.1 PPT
Hydrogen Maser	Passive	1 PPT
Hydrogen Maser	Active	.0007 PPT

When you talk about parts in a trillion the choices become somewhat limited although they are more readily available today to the experimenter than a few years ago. Two excellent reference sources have become available in recent years.

1. Rubidium standards are appearing as surplus items because they are being retired from service in cellular telephone systems around the world.
    1. Aging reduces their output, but because the frequency is determined by a physical property, their accuracy remains in tact.
    2. Frequency-stable outputs can be accurate at the parts per trillion ($10x10^{-12}$) level.
2. Crystal oscillators can be “disciplined”--phase-locked--to U.S. Global Positioning System (GPS) satellites.
    1. The GPS system requires that individual satellite clocks be maintained at accuracies exceeding a few parts per trillion.
    2. Earth-bound oscillators can be synchronized with the GPS satellite clocks, using timing information derived from the GPS signal.

3. Calibration

3.1. Measurement

Given the availability of such references, the *metrologist* performing a verification or calibration must answer two important questions:

A. Is a single, brief frequency comparison sufficient to ensure an accurate measurement?

B. In making the measurement, how does one compare two frequency sources and estimate the difference/error between them?

The answer to the first question is No.

A. Separate measurements an hour, a day, or a week apart do not confirm that the difference between a reference and a Device Under Test (a DUT) is consistent.

B. At best, they can only confirm that the measured error was the same at two points in time. They could not show that the DUT frequency remained constant through the measurement period.

C. Many measurements at short intervals over a long time are needed to establish the variation of the DUT with time.

1. The results can be plotted to show, graphically, how the reference and DUT frequencies differed over time.
2. The results can be “averaged” to arrive at a crude figure of merit for the DUT.

A number of answers can be given to the second question:

A. A frequency counter with a sufficiently accurate time base can be programmed to accumulate measurements over a period of tens, hundreds, or even thousands of seconds. The long measurement time base improves overall accuracy by integrating errors over time; however transient deviations are averaged.

B. Relative phase method - when the frequencies of the reference and the DUT are nearly the same, measuring the difference in the phases of the two signals may also be possible.

1. Oscilloscope - With both a reference and DUT signals applied to the inputs of an oscilloscope, the scope can be triggered by say the reference signal and the relative phase can be compared at time A and at time B. Then with the difference in time and the difference in phase either in degrees or in time the difference in frequency can be calculated. An example is shown in Figure 2. Here the time is 240 seconds and the phase is 156 - 170 or 14 degrees to yield a difference in frequency of 16 PPT. This example is rather straight forward in that this oscilloscope computes phase. An eyeball estimation of the phase is somewhat tedious.

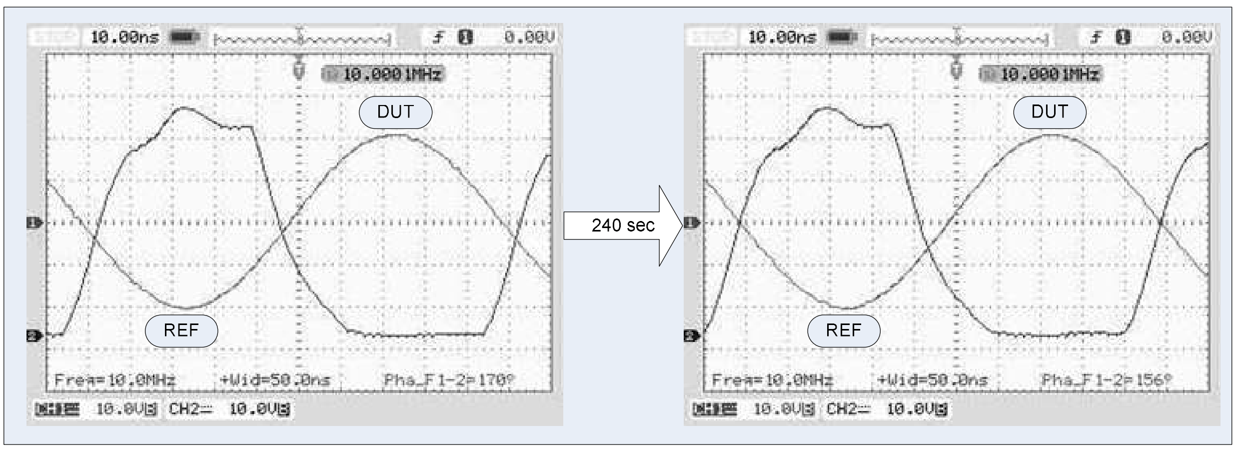

Figure 2 Example Relative Phase Measurement with Oscilloscope

2. Stripchart - My first attempt at verifying the accuracy of my HP106B standard involved comparing its output to WWV at 60 kHz and recording the relative phase of the WWV reference and the DUT on a stripchart recorder.

a) This method can be time consuming because resolving a difference of only 10 degrees at an accuracy of $1x10^{-11}$ would require 12.9 hours. Increasing the resolution accuracy to $1x10^{-12}$ would require extending the measurement time to 5.33 days.

b) The measurement time can be reduced when tests are done at higher frequencies.

For 60 KHz, a frequency error of 10 PPT represents a phase change difference between the reference and the DUT of 0.0002156 degrees per second. Thus to accumulate a phase change difference of 10 degrees would require 46382 seconds or 12.9 hours. And, for a frequency error of 1 PPT would require 463820 seconds or 5.33 days.

For 10 MHz the times are much shorter. for the same error levels as for 60 KHz the times would be 278 and 2783 seconds respectively.

Of course, being resolving the frequency difference change accurately depends on being able to resolve the phase difference change from the medium presented such as the stripchart or oscilloscope accurately. That is another matter entirely.

3.2. Adjustment

Having measured the frequency difference the next task in calibration is to adjust the DUT to some given limit or as close as possible to the reference frequency. When long measurement times are encountered the adjustment process becomes problematic since the operator cannot see the effect of the adjustment until another measurement is made, It takes time and patience to make this process converge. If the operator can see the effect in real time the process becomes much easier and straight forward.

4. Relative Phase Method

The Relative Phase Method for comparing the frequencies of two sources can be extremely accurate if an accurate reference is available.

1. The method determines the rate of change of phase between two sources. A computer monitors the phase difference between them and presents the error graphically.
2. For example, a one-Hertz phase error represents a phase change of 360° per second. At 10 MHz, the time required to complete one Hertz (one RF cycle) is 100 ns, so the error due to phase change can be described as 100 ns per second.
3. We can elaborate the idea by stating the change in degrees of phase error per second for a given fraction of a cycle. Thus for a 10 MHz comparison:

Table 2 10 MHz Error Values vs. Frequency Error

Error - cycles/sec	PPM/PPT	Deg/sec	Nsec/sec
1	0.1 PPM	360	100
0.1	0.01 PPM	36	10
0.01	0.001 PPM	3.6	1
0.001	100 PPT	0.36	0.1
0.0001	10 PPT	0.036	0.01
0.00001	1 PPT	0.0036	0.001

4. Clearly, small phase errors approaching 0.0036 degrees/second or 0.001 ns/second become visible only after many seconds if they are to be observed on an oscilloscope. A 1 ns error, for example, completes a full 360° phase shift in 1×10^9 seconds, which is nearly 280 hours, or 11½ days.
5. An important human factor involved with these types of measurements on an oscilloscope is keeping track of the variations over long periods of time.

From Table 2 some useful conversion constants emerge.

- deg/sec = (1/.0036) PPT
- deg/sec = (1/360) PPM
- Nsec/sec = 1000 PPT
- Nsec/sec = .01 PPM

This method must be automated if accuracies approaching 1×10^{-12} are to be realized.

5. Monitoring

Something I have found interesting in working with precision oscillators is that rarely are they what they seem at first glance. Since I have designed and built the instrument to be described here I have found the true performance of several standards I have in my lab to be less ideal than I had expected. Because you calibrate an oscillator to 1.1×10^{-10} and you come back tomorrow and it reads 1.2×10^{-10} doesn't necessarily mean that you can depend on say that it has 2×10^{-10} accuracy for general measurements. Now that I have the capability of monitoring various sources at the 10^{-11} or 10^{-12} level, I have greater reservations about the performance I can expect from a given source with a certain calibration level. If you want to know how a source will truly perform, I believe it is essential to monitor the error performance over a period of time and from that determine what you can count on for the performance of that source.

6. The Test Set

I chose to accurately measure the relative error between a reference standard of known accuracy and an unknown device under test (DUT) dynamically in real time. I determined that the best candidate for this measurement was an accurate phase measurement and mechanization of the method with as much precision as I could bring to bear on the implementation. Further, I wanted to present the results in real time on a computer screen for observation. In this way I could observe the real time performance of various standards. What follows is the result of the evolution that development.

6.1. Overview

The basic concept is simple as shown in Figure 3. The specifics get a little more interesting.

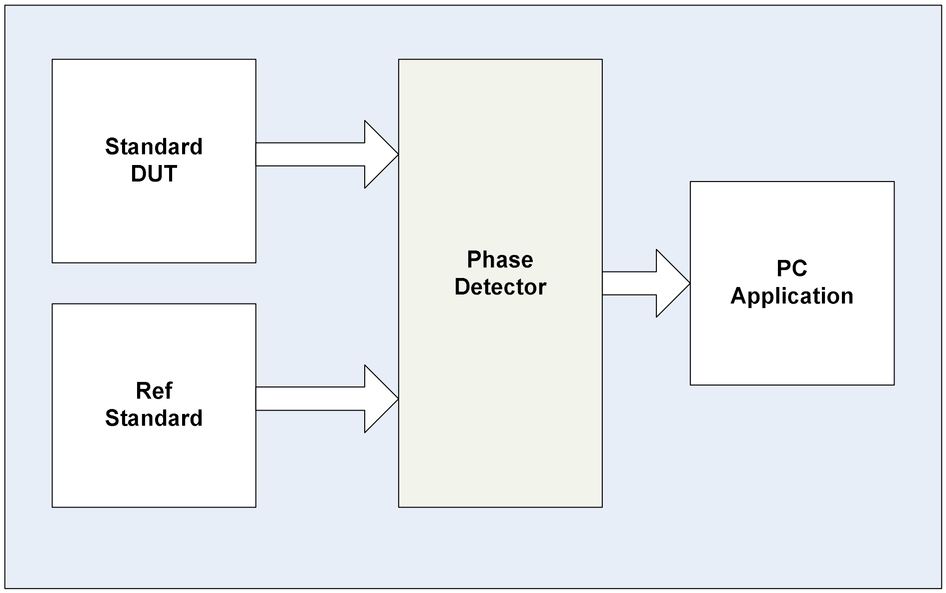

Figure 3 Basic Concept

The figure shows that the phase of the reference and DUT are compared and a value relating to the phase difference is passed to the computer. The computer supplies the time base for the calculations that result in the resultant error being displayed on the computer screen.

Figure 4 shows the detailed block diagram of the test set. To achieve high accuracy the signals must be conditioned significantly. Each channel is first detected and phase locked to a higher frequency crystal VCO to obtain a precise square wave for presentation to the phase detector. The reason for this is that in simply squaring the input signals, it is difficult maintain the necessary 50% duty cycle required for precise phase determination. In addition, for the quadrature phase detector chosen the reference frequency must be four times the reference frequency. This will be discussed below.

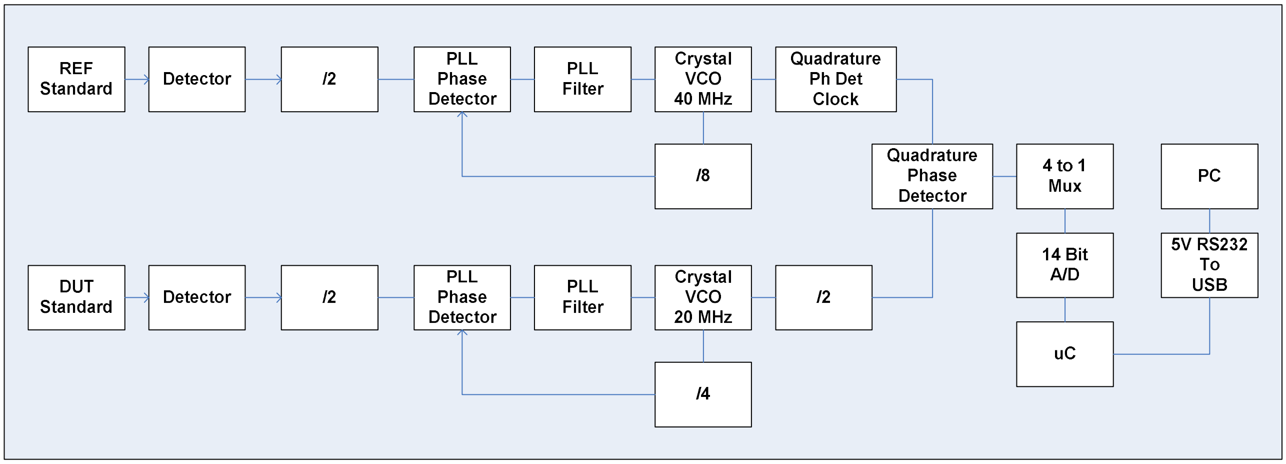

Figure 4 Detailed Block Diagram

The Reference and DUT channels are almost identical. The exceptions are that the Reference channel operates at 40 MHz and its output supplies clocks for the quadrature phase detector. I will discuss the quadrature phase detector below.

The Reference input signal is converted from analog to digital by the input detector and then is divided by two by a digital counter. The 5 MHz output feeds the 74HC4046 PLL phase detector. I prefer to operate this PLL phase detector at somewhat less than 10 MHz. The output of the PLL phase detector feeds the crystal VCO via the PLL filter. The 40 MHz VCO output feeds a divide by eight to feed 5 MHz to the PLL phase detector to close the 40 MHz PLL loop and it drives a circuit to derive the reference clocks for the quadrature phase detector.

The DUT channel is identical to the reference channel except that the crystal VCO is operated at 20 MHz and the output is divided by two to feed the quadrature phase detector at 10 MHz and divided by four to feed the PLL phase detector at 5 MHz and close the 5 MHz PLL loop.

The outputs of the Reference and DUT channels are compared in the Quadrature Phase Detector (QPD) and fed to a 14 bit A/D that is coupled to the microcontroller via the Serial Peripheral Interface (SPI). The microcontroller takes in the digitized phase values and routes them to the Personal Computer (PC) via a 5V RS232 to USB adaptor. the PC running a Windows application processes the data and presents the analyzed results on the PC monitor. The PC application is written in National Instruments Lab Windows/CVI.

6.2. Quadrature Phase Detector

The phase detector chosen for this implementation is the exclusive OR logic gate (XOR). For a single XOR the output waveform for a phase difference from 0 to 360 degrees is shown in Figure 5. The output when filtered is a triangular wave in amplitude vs. phase difference. Note that the output is multi-valued, that is there are two phase values for each voltage value, therefore the unique phase cannot be determined from the voltage reading without additional information.

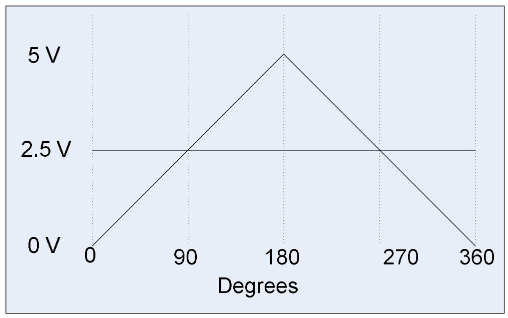

Figure 5 Single XOR Phase Detector Output

This can be remedied by the addition of a second XOR driven by the reference 90 degrees out of phase forming the QPD. From the reference clock multiplied by four the QPD reference signals are derived. The result is shown in Figure 6and overlaid on the same trace in Figure 7.

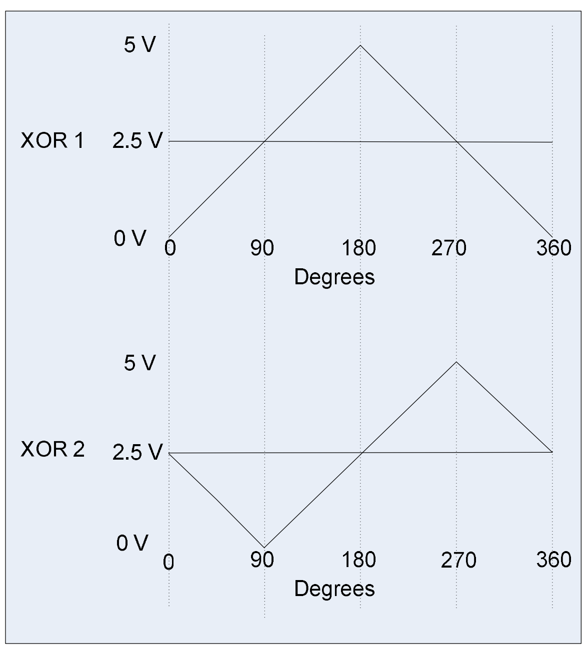

Figure 6 Quadrature XOR Phase Detector Outputs

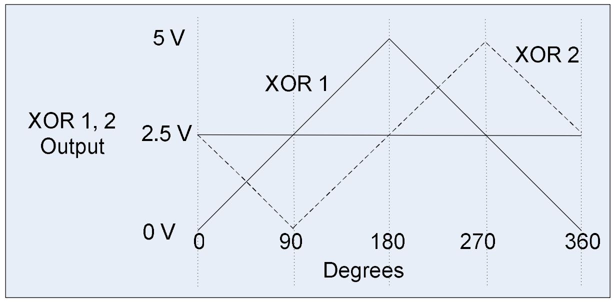

Figure 7 Quadrature XOR Phase Detector Outputs - Overlay

From the combination of the two outputs the in-phase Component I and the quadrature component Q the unique phase can be determined, since there is a unique value pair for each phase point.

The circuit I use for the generation of the quadrature clocks is shown in Figure 13. I use an additional two XORs to help with a very slight nonlinearity of phase vs. amplitude at the peaks and valleys of the output. For this I generate 10 MHz clocks at 0 (I), 90 (Q), 180 (-I) and 270 (-Q) degrees. Figure 8 shows the XOR outputs using the four clocks.

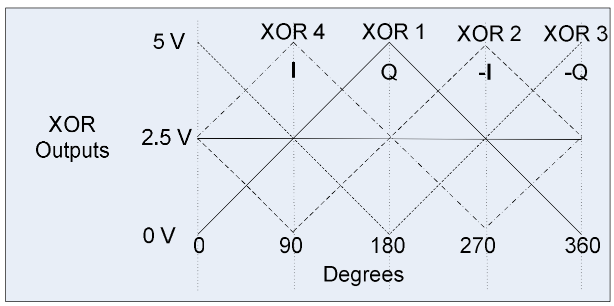

Figure 8 Outputs of the Four XORs vs. DUT Phase

The I and -I and the Q and -Q are combined in software to give the resultant I and Q from which the phase calculations are made.

Figure 9 shows a screen presentation from the GUI for a period of 98.5 seconds using a Trimble Thunderbolt GPS supervised source as the reference and an ovenized crystal oscillator as the DUT. There are two stripchart presentations. The lower stripchart is the measured data showing the measured I and Q and the calculated phase. From the calculated phase the short and long term errors along with the accumulated phase are calculated and presented in the upper stripchart in the figure. In the GUI colors are used to help identify the various traces. Here the long term error is 4.01 PPT over a period of 98.5 seconds and the last instance of short term error is 3.48 PPT.

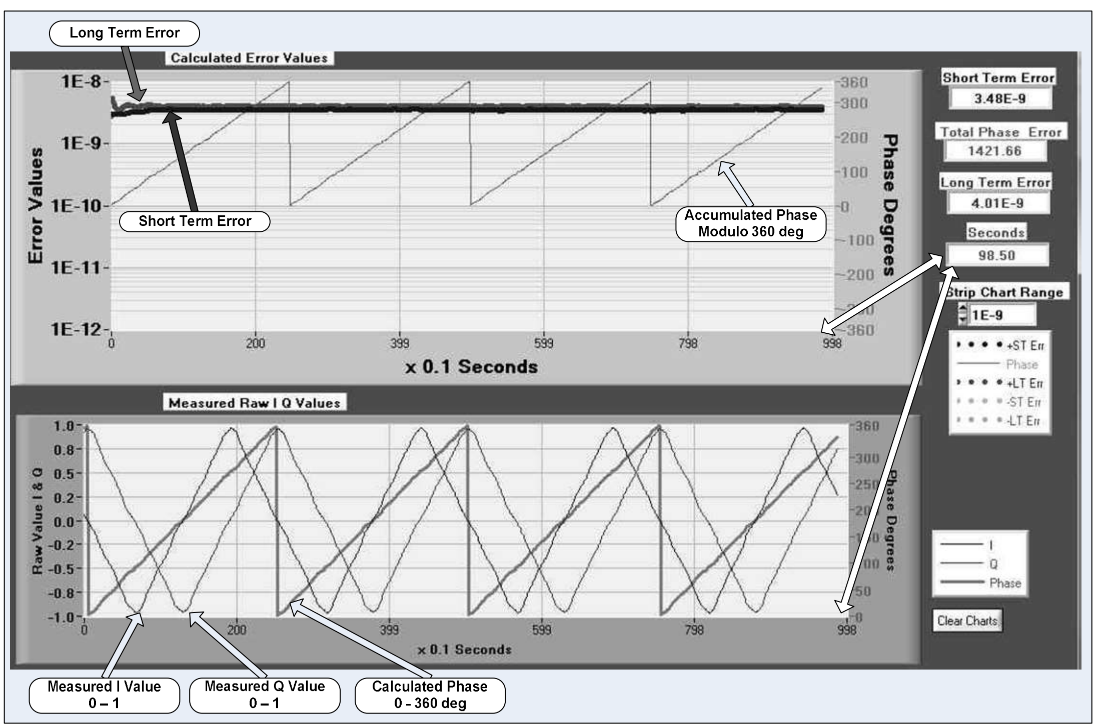

Figure 9 Example Quadrature Phase Detector Output

Figure 10 shows a annotated screen presentation of the QPD using a Trimble Thunderbolt GPS supervised source as the reference and a rubidium oscillator as the DUT. Here the phase movement is so slow that QPD outputs look almost stationary. The top stripchart shows the error values and the accumulated phase difference. I beg the readers indulgence in viewing this stripchart as color on the GUI makes the delineation of traces easier.

On the GUI the red/pink trace is the short term error and the blue/light blue trace is the total accumulated averaged error since the test began. Since negative values are problematic on a log chart, the error values are differentiated for sign by color shading with the dark colors representing positive error and the light shading being the negative error. The green trace is a result of the accumulated phase in degrees. For this display the phase display is limited between -360 and 360 degrees and we see the trend in phase with time.

Note that the scale on the left has been adjusted for this condition. If one studies the traces closely one can determine the extent of the long term error trace as it stabilizes and remains rather convergent while the short term trace varies considerably. This test is approaching the limit of the test set as can be seen below in the section on performance. Note that the total phase error for the 98.8 seconds is only -0.54 degrees giving a long term error of - 1.52 PPT.

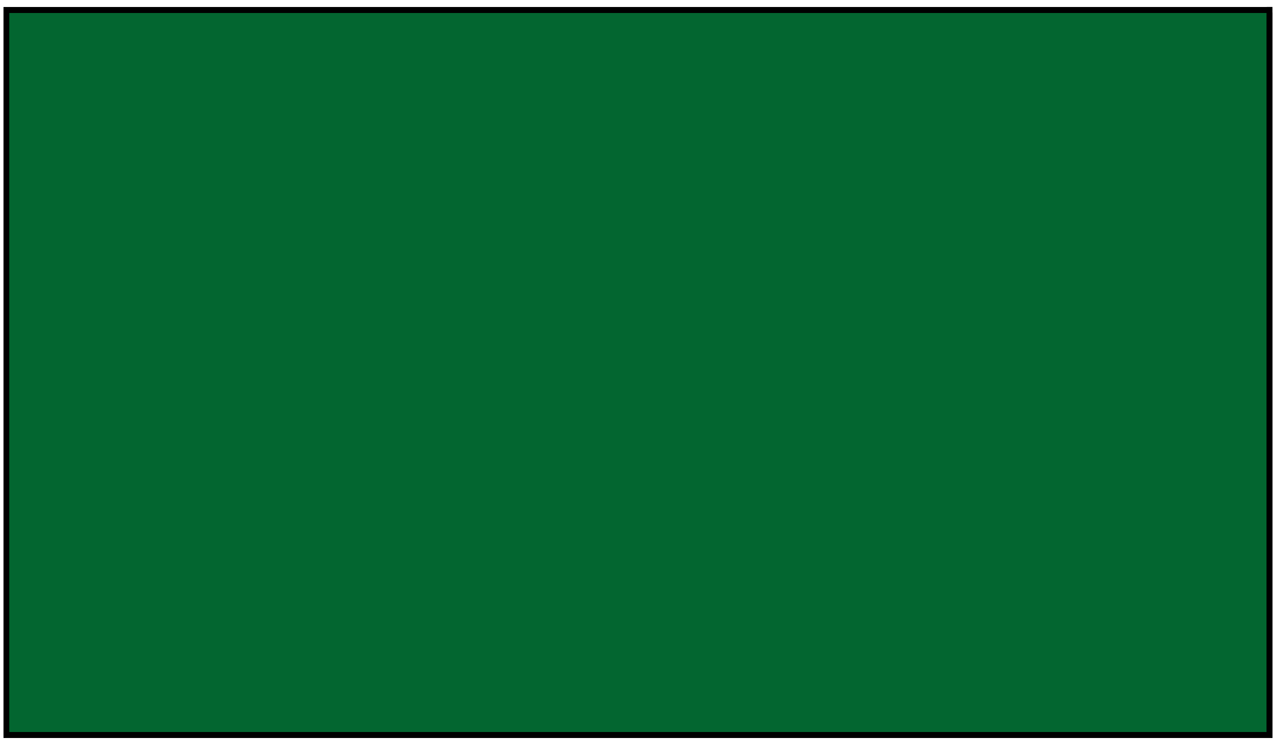

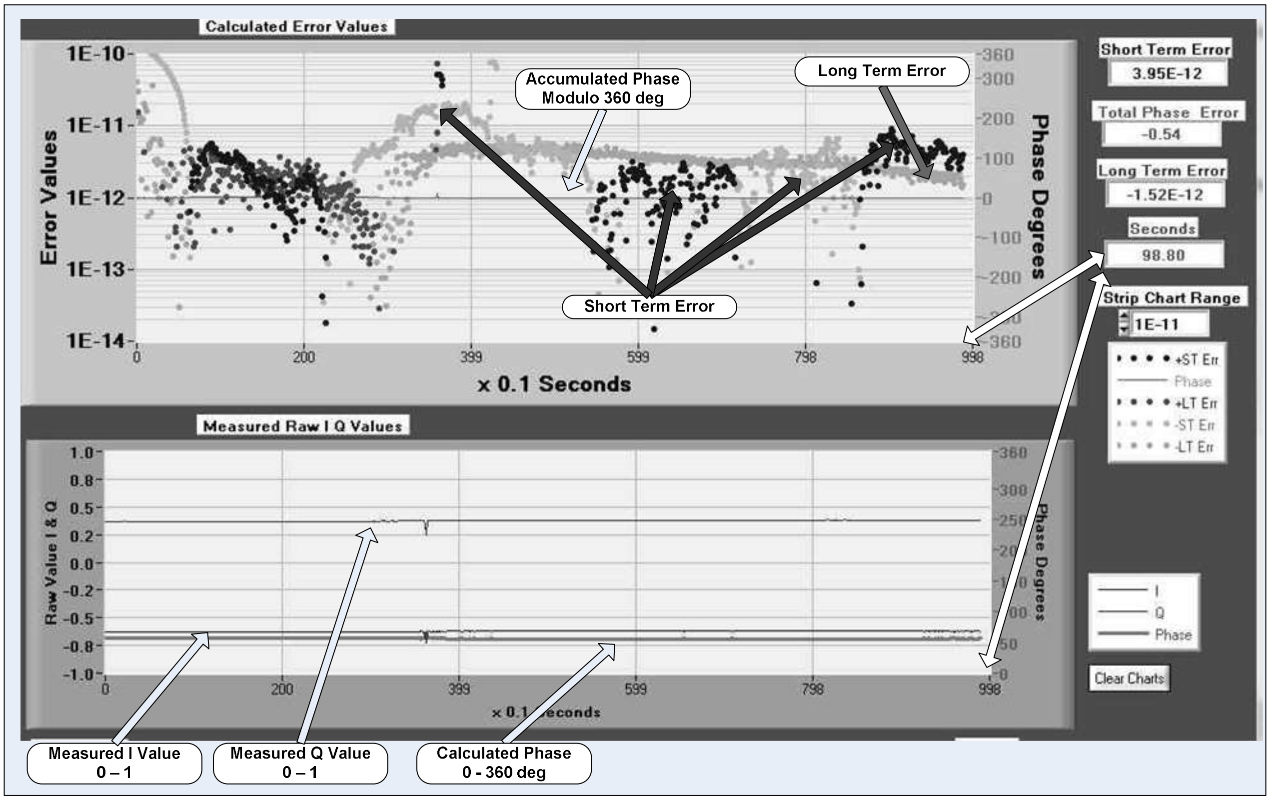

Figure 10 Example Output Data

6.3. High Level Circuit Description

The input detector (Figure 11) is designed to accept a variety of signal input levels and types.

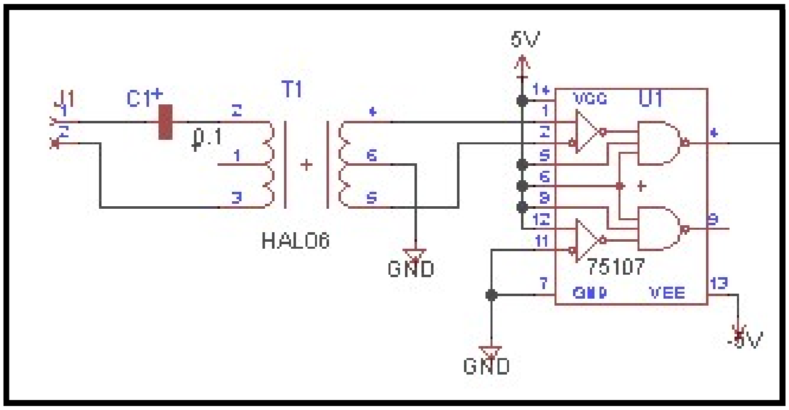

Figure 11 Input Detector

The transformer provides ground isolation between the two sources. The 75107 line receiver provides a reliable transition detector function.

The VCO does not necessarily give a symmetric waveform output, so a divide by two is necessary to guarantee essential symmetry. Symmetry is particularly important in driving the phase detectors to derive phase as precisely as possible.

The quadrature phase detector shown in Figure 12 offers the ability to uniquely identify the phase difference from 0 to 360 degrees. Only two XORs are necessary to accomplish this; however the second two allow a bit more resolution as will be discussed below.

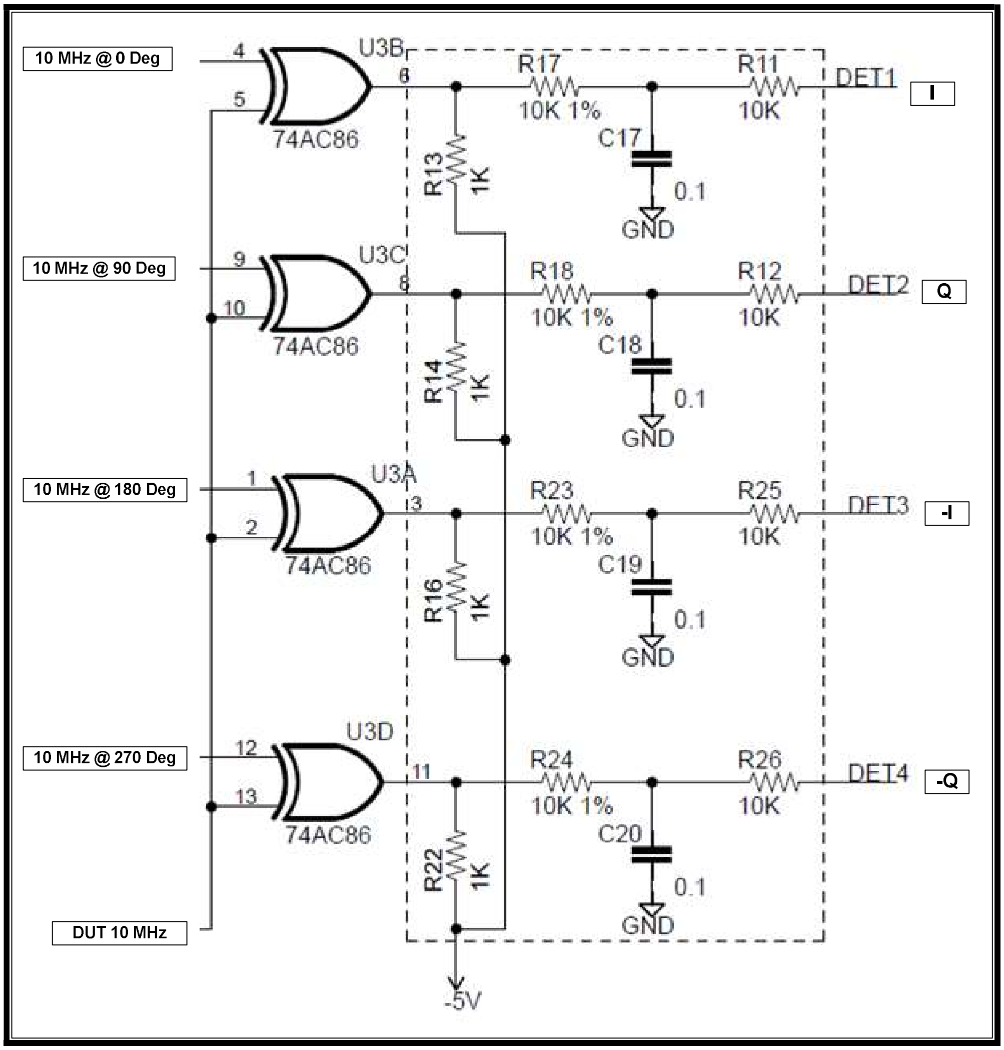

Figure 12 Quadrature Phase Detector

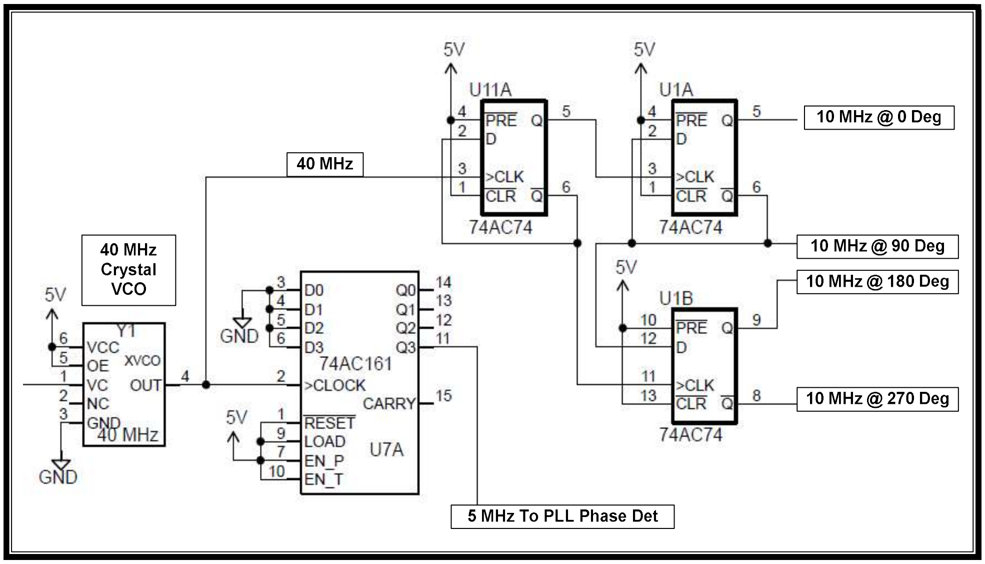

Figure 13 Quadrature Clock Circuit

To insure symmetry the circuit must be driven by four times the output frequency.

7. Hardware

The hardware is constructed using breadboard techniques and thru hole DIP ICs. The microcontroller is an Atmel ATMEGA328P module.

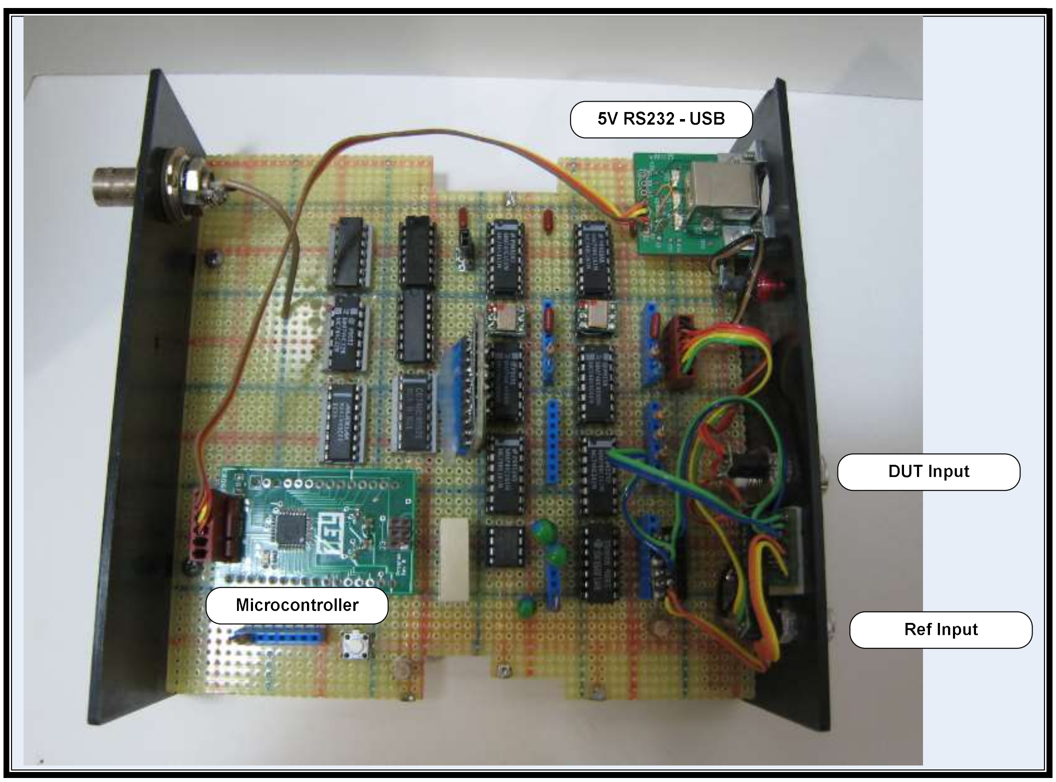

Figure 14 Test Set Component View

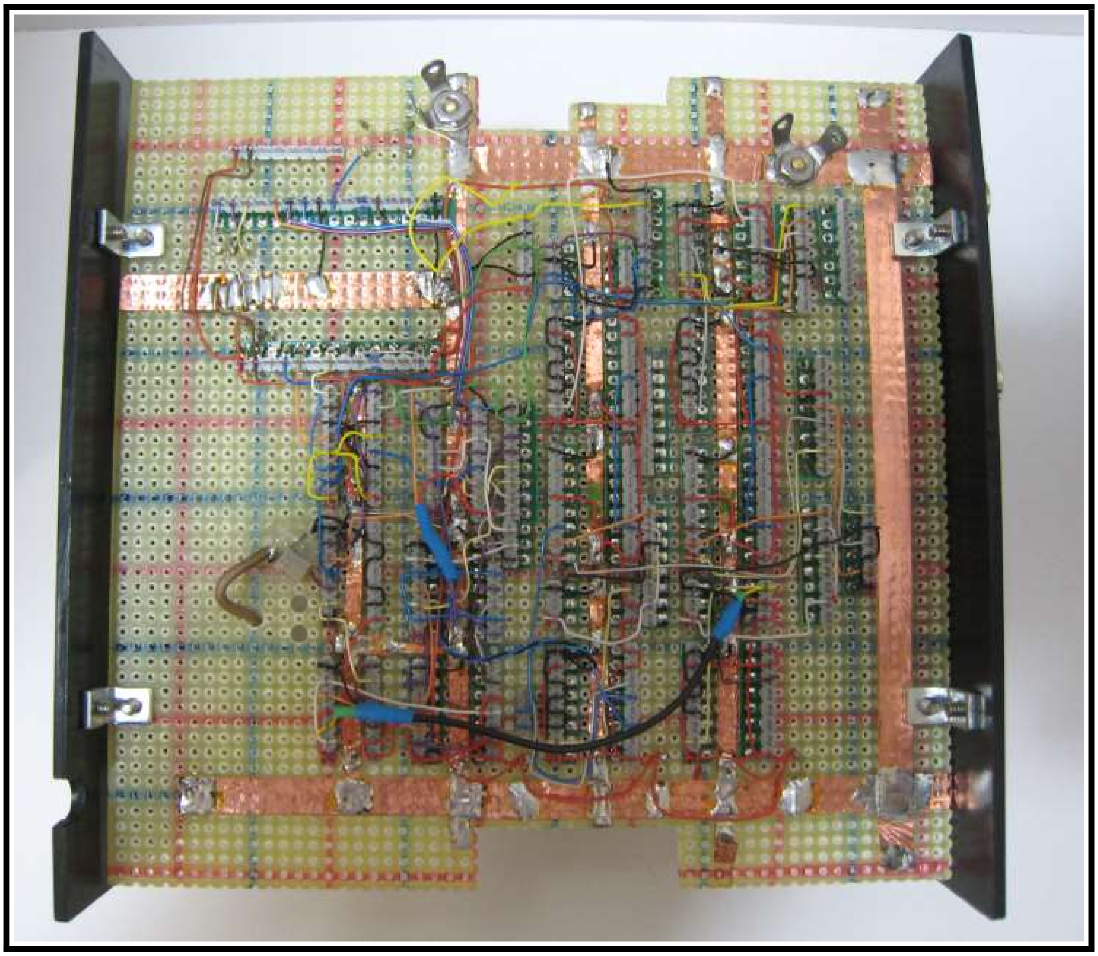

Figure 15 Test Set Wiring View

Figure 16 Test Set Front

8. Software

8.1. Microcontroller Firmware

The microcontroller code is written in C for the Atmel ATMEGA 328P.

8.2. Windows Application

The Windows application is written in C using National Instruments Lab Windows/CVI.

9. Performance

It is interesting to look at the parameters that limit the performance of the test set. Of course the primary limit is the accuracy and stability of the reference source. Beyond that is the accuracy and stability of the PLLs. Further the short term accuracy is particularly dependent on the resolution of the A/D.

9.1. Reference Sources

It is important to note that for a meaningful reading the reference accuracy must be at least an order of magnitude better than the desired level of measurement. There are a number high accuracy reference sources; however cost becomes a big factor for me. Recently two sources have become available at economic prices on the used market. They are the Trimble Thunderbolt GPS supervised source and various Rubidium sources.

9.1.1. Trimble Thunderbolt

There are at least two kinds of GPS supervised sources. The first is that most generally seen in experimenter circles that uses the GPS one-pulse-per-second (1PPS) to compare with an ovenized crystal oscillator or a Rubidium source and derive a correction signal to correct the source. The second and the technique that Trimble uses is to derive the GPS receiver reference clock from an ovenized crystal oscillator, then use the timing error from the GPS position solution to correct the ovenized crystal oscillator.

It is interesting that Trimble does not specify the accuracy of the Thunderbolt 10 MHz output; however they do specify the phase noise. I suspect the that this is due to the various potent user related situations where a solid GPS solution is not possible. Such situations might include the setup or the antenna placement. The Windows utility they supply to setup and monitor the unit in operation does give a dynamic error report and with the unit I have that error varies around in parts in 10^{-12} . Some thought about the stability of these types of sources reveals that at some level they must jitter around. Consider that the source is an ovenized crystal oscillator and with a very few exceptions it is very difficult to make them perform better than $1x10^{-10}$. To me that means that in this application with the oscillator controlled to $Xx10^{-12}$ it is constantly moving away from that point and must be constantly brought back by the controller. The exceptions mentioned refer to ovenized crystal oscillators built as instruments such as the HP 106B that can perform to a higher level. For my work I consider the Thunderbolt to be able to hold something like $5x10^{-12}$.

9.1.2. Rubidium Sources

A number of Rubidium sources have come available on the used market as the result of cellular system upgrades. I have read and heard mentioned that Rubidium sources age out and become useless; however in one of the Datum[1] papers it states that this is not true.

9.2. Phase Locked Loops

The PLLs must very, very closely follow the reference and DUT sources for an accurate short term measurement. In a PLL the VCO output is constantly compared to the reference frequency and an error correction signal constantly controls the VCO. Inherently the error correction signal will have some noise and there will also be some error. If the PLL is locked, that only means the VCO phase stays in a given region relative to the reference. There is always some jitter or variation. The objective here is keep that jitter very, very small.

Long term measurements are less critical because, although the PLL VCO may jitter about the input frequency it is still locked and the longer term phase difference is still valid even though the short term phase difference may be significantly in error.

9.3. A/D Converter

The resolution of the A/D converter is very important. The A/D employed is 14 bits thus giving the LSB to be one part in 16384. Since the range of one slope of the QPD output is 90 degrees, one LSB represents 0.0055 degrees. To improve this resolution a number of readings are takenfor each point and averaged to yield the value used in the calculations.

9.4. Example Encompassing All Factors

As is sometimes said the proof of the pudding is in the eating. All of the factors limiting the range of the test set come into play when both the reference and DUT inputs are fed from the same rather highly performing source. Figure 17 shows the result of feeding the test set with the same source on both the reference and DUT inputs. For the result shown the two PLLs have to be operating independently in the vicinity of $1\text{x}10^{-12}$. Otherwise, the phase difference between the two channels would be much greater. One factor that does not come into consideration here is that the two are operating from the same frequency and therefore are not moving relative to each other. I have observed in earlier layouts interaction between the two channels when the zero crossings of the two channels are in the same region in time. Notice in this case that the long term error is predominately positive and that it begins to settle out as time progresses while the short term error continues to vary about. Again I beg the readers indulgence in interpreting the presented data.

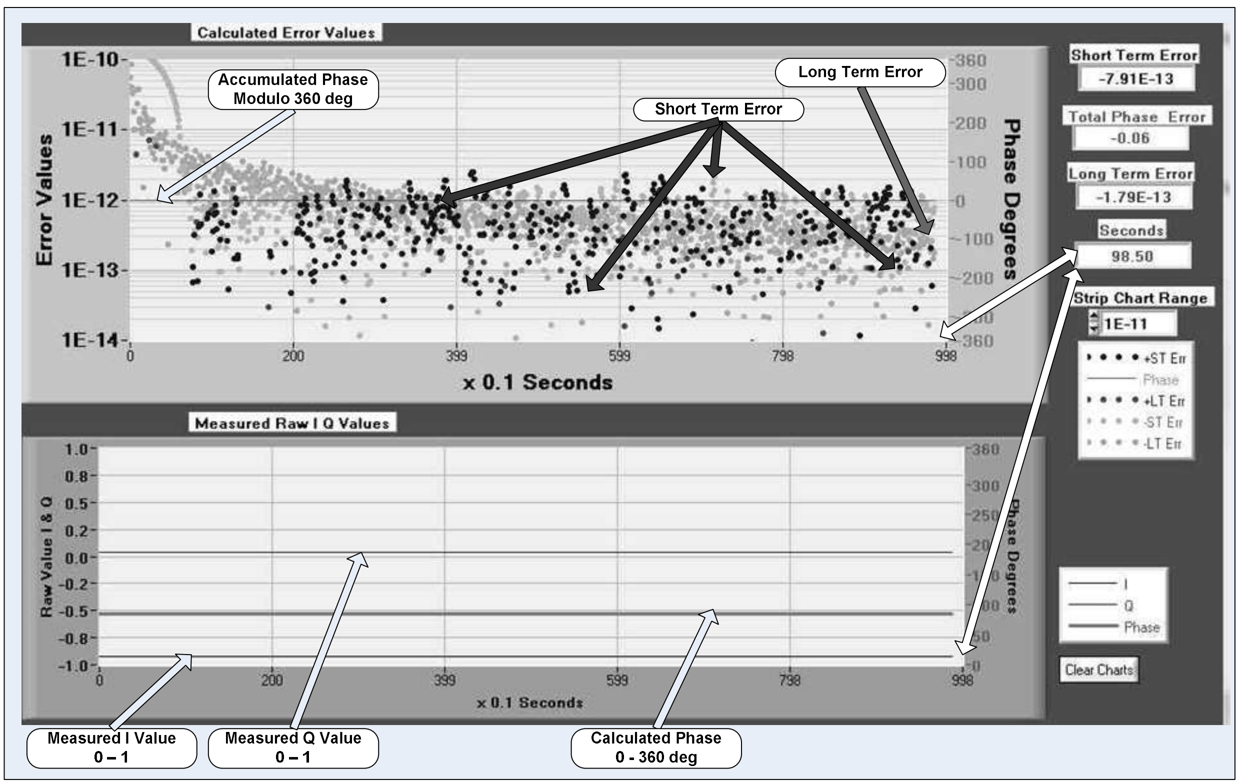

Figure 17 Same Source (Trimble Thunderbolt) to Ref & DUT Inputs - Each Time Inc 0.1 sec

10. Schematic

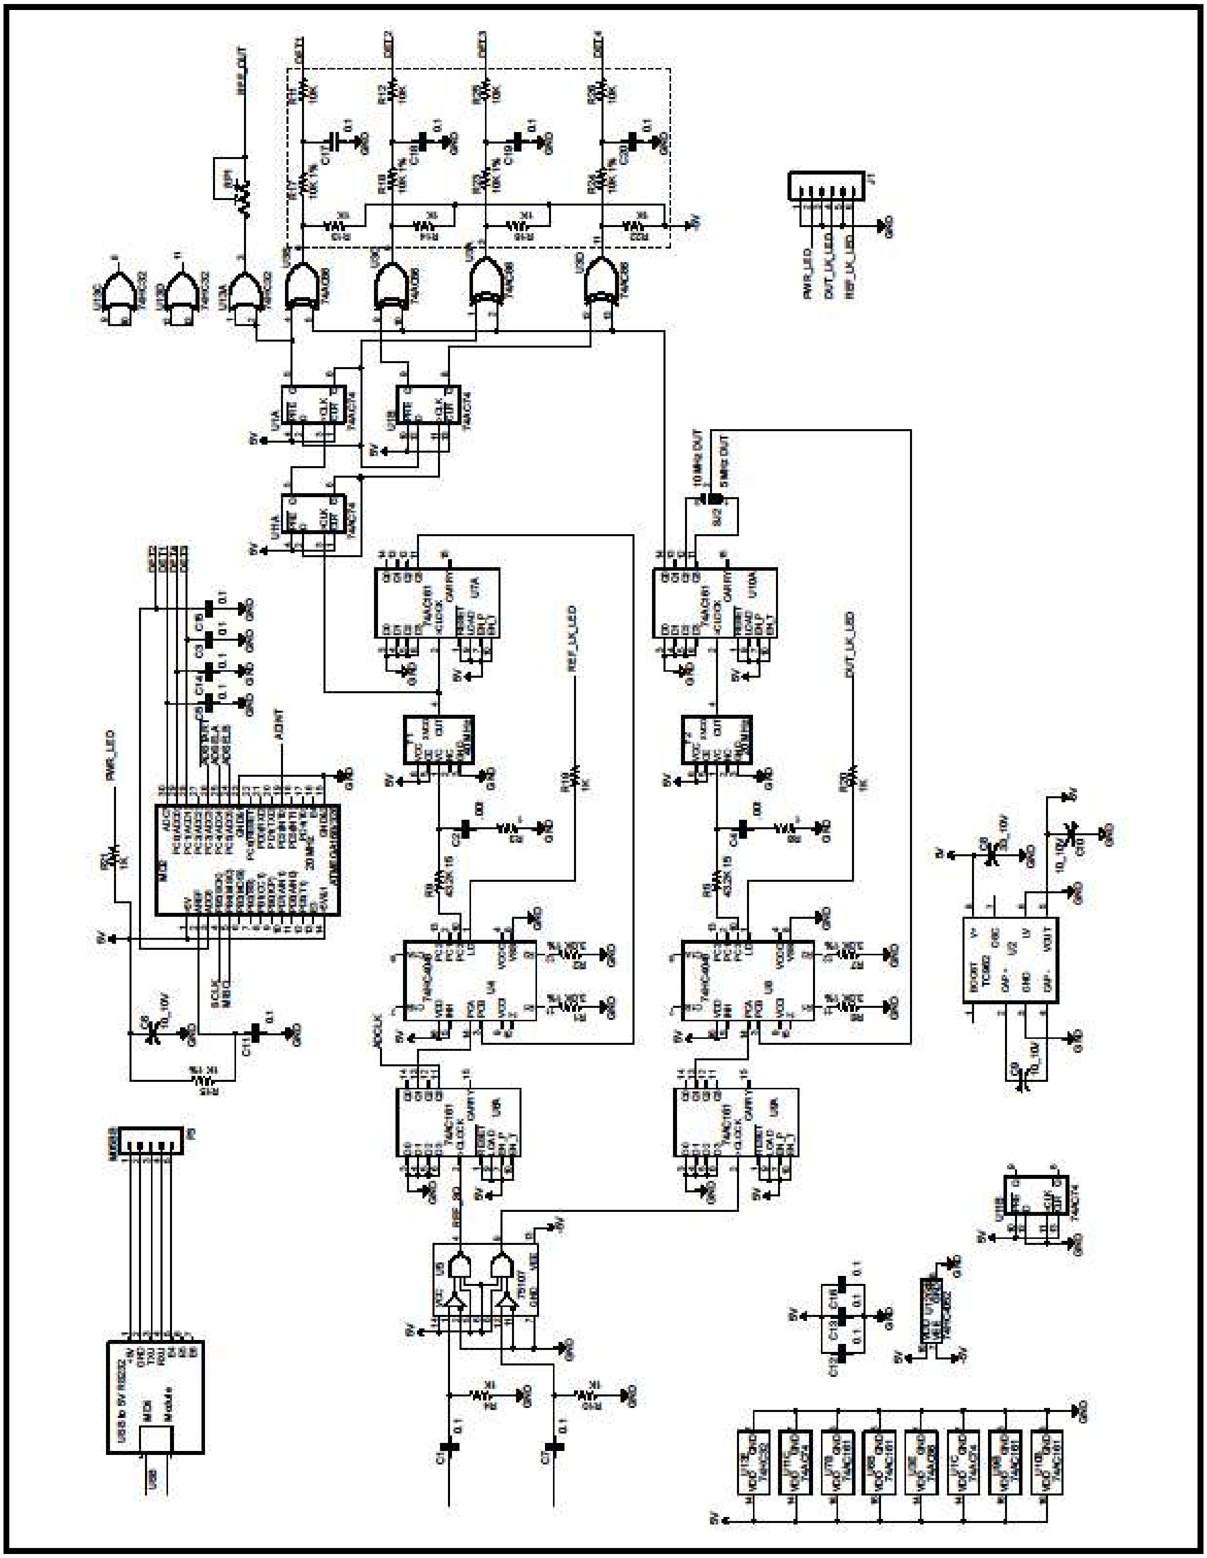

Figure 18 Schematic Sheet 1

Figure 19 Schematic Sheet 2

11. References

1. "A Note about Rubidium Oscillators", Datum, 12/10/1998,
http://www.amtestpl.com/DATUM/GmbH/Rubidium-Note-e.pdf

1. The Stable32 Program for Frequency Stability Analysis, Hamilton Technical Services, Beaufort, SC
29907, http://www.wriley.com.
2. D.B Sullivan, D.W Allan, D.A Howe, and F.L Walls, eds., "Characterization of clocks and oscillators," Natl. Inst. Stand. Technol. Technical Note 1337, **http://tf.nist.gov/timefreq/general/pdf/868.pdf** (March 1990).
3. D.A. Howe, D.W. Allan, and J.A. Barnes, "Properties of signal sources and measurement methods," *Proc.*
35th Freq. Cont. Symp., pp. 1-47,
http://tf.nist.gov/timefreq/general/pdf/554
.pdf (May 1981).

NIST Special Publication 1065

A Tutorial Introduction" available here.

http://tf.nist.gov/timefreq/general/pdf/868.pdf

http://tf.nist.gov/timefreq/general/pdf/554.pdf

End of Document

L:\JimData\!\!FreqPhase\Docs\FreqPhaseCal0_08.doc