

Copyright 2002 Agilent Technologies, Inc. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works, must be obtained from the Agilent Technologies contact:

Dr. William R. Shreve,  
Laboratory Director, Systems and Solutions Laboratory  
Agilent Technologies  
3500 Deer Creek Rd., MS 24M-A  
Palo Alto, CA 94304-1392, USA  
650-485-2664

By downloading or reproducing this document: (1) you acknowledge that Agilent Technologies, Inc does not warrant the accuracy or completeness of the material and disclaims any responsibility for it and (2) you agree to include this complete paragraph in each copy that you make.

Agilent Laboratories  
3500 Deer Creek Rd.  
Palo Alto, CA 94303

Date: 25 November 2002



## Node Synchronization Improvement by High Stability Oscillators

Michael C. Fischer

Fischer Consulting, 763 E. Charleston Rd., Palo Alto, CA, 94303  
650-855-9816 (phone), [mikecf@jps.net](mailto:mikecf@jps.net)

John C. Eidson

Agilent Laboratories, 3500 Deer Creek Rd., Palo Alto, CA, 94304  
650-485-4263 (phone), [john\\_eidson@agilent.com](mailto:john_eidson@agilent.com)

Bruce Hamilton

Agilent Laboratories, 3500 Deer Creek Rd., Palo Alto, CA, 94304  
650-485-2818 (phone), [bruce\\_hamilton@agilent.com](mailto:bruce_hamilton@agilent.com)

### ABSTRACT

Distributed measurement and control systems in which all components share a common sense of time allow significant decoupling of synchronization issues from considerations of communication latency and latency fluctuation. This decoupling enables new techniques for solving measurement and control problems involving large numbers of sensors and actuators or complex synchronization requirements. A common sense of time is established by having each component contain a real-time clock synchronized to its peers. This may be accomplished using IEEE 1588, a standard that defines a protocol enabling the precise synchronization of clocks in the components of a networked, distributed measurement and control system. Important applications of this technology could benefit from improvements beyond the basic accuracy exhibited by the initial implementations of the standard. Synchronization accuracy improvements obtained by using a higher stability quartz oscillator and modified tracking loop parameters are tested and results reported.

### INTRODUCTION

Most of the systems supporting command, control, communication, identification, navigation and other vital functions either rely critically or are aided by reference to an accurate sense of time that is common to all elements of a particular system and to other systems that interoperate. The distribution of such a coordinated time reference over a communication network is a feature of the Network Time Protocol (NTP) via Ethernet. However the accuracy of NTP is on the order of a few milliseconds. A recently established standard, IEEE 1588, holds promise for achieving sub-microsecond synchronization of real-time clocks over a localized area. Such areas can be coordinated by providing each area with a master clock synchronized by GPS, for example. Initial prototype implementations of this standard have verified the expected levels of performance, and analysis suggests that the instability of the low cost quartz clock oscillators employed might be the source of much of the remaining instability. This concept led to a project to replace the low cost oscillators with the best quartz oscillators available and measure the effects on performance. The results of the project are presented in this report.

## **REQUIREMENTS FOR MILITARY APPLICATIONS**

While military applications were not explicitly considered during the development of IEEE-1588 it appears that many of the characteristics of industrial and test and measurement systems are applicable to the newer generations of military systems. In particular military systems are evolving from stand-alone systems into an architecture of interoperable systems with strong synchronization or coordination requirements. A common sense of time is to be the foundation for this architecture with the worldwide GPS system as the key component. However operation must be insensitive to the loss of GPS signals, [6]. At least in local environments such as a ship, or command center a system based on IEEE-1588 may allow more robust local operation by maintaining local time consistency even when isolated from the global system. As system requirements extend to smaller, less capable and cheaper military devices the low implementation footprints for IEEE-1588 may be advantageous as well.

Further discussion of applications of IEEE 1588 may be found other papers, [1], [5]. The remainder of this paper discusses performance experience and how IEEE-1588 is structured to meet timing requirements.

## **IEEE-1588 DESIGN FOR MEETING TEMPORAL REQUIREMENTS**

Within a subnet IEEE-1588 automatically establishes a master-slave relationship among the participating clocks communicating via the subnet. The master is selected as the best clock based on defined descriptors for inherent accuracy, traceability to UTC, etc. Provision is made to designate a set of clocks to be preferred in this selection for applications where this is important.

The slaves synchronize their local clocks to that of their master by an exchange of messages illustrated in Figure 1. Periodically the master clock sends a distinguished message, a Sync message, as a multicast to all its slaves. The master includes an estimate

of when the Sync message will be placed on the network. In the most accurate implementations the master will contain a mechanism for detecting and time stamping based on the master's local clock, the time that the Sync message is actually placed on the network. In Ethernet an ideal place to attach this detector is at the MII interface of the PHY chip, thus avoiding the large temporal fluctuations in the upper portions of the protocol stack. If the master is so equipped, the master's IEEE-1588 code sends this measured time stamp to all slaves in a second message, the Follow\_up message. The slaves receive the Sync message and if so equipped detect and time stamp its arrival as close to the network as possible. Upon receiving the Sync message and for highest accuracy the Follow\_up message and the local receipt time for the Sync message, the slave's IEEE-1588 code uses this information to correct the time of the slave's local clock. Periodically, but with longer period to reduce network loading this process is reversed. This forward and reverse path information is used to compute the one-way network latency on the assumption that the path is symmetrical. The slaves use this measured latency in computing the correction to their local clock. This procedure effectively removes timing fluctuations in the participating end devices and latency in the communication path.

The other major source of timing errors is network latency fluctuations introduced by network elements. In an Ethernet environment this includes repeaters, switches, and routers. Of these, routers introduce fluctuations too large and inconsistent to be reduced to the desired accuracy with statistics. IEEE-1588 specifies a transfer standard mechanism, the boundary clock, for logically eliminating routers from the IEEE-1588 protocol communication path. The boundary clock appears to each subnet of interest as an ordinary IEEE-1588 clock as if it were an end device. However the boundary clock uses a single local clock to serve as a master in all or all but one of the subnets of interest. In the rare case where multiple routers are needed in an application, the boundary clocks themselves form a hierarchy so that in a properly implemented IEEE-1588 system there is always a single clock serving as the primary source of time for the ensemble. The time base of this clock may be synchronized to a reliable source of UTC time if required.

In the restricted environments typical of the target applications, the fluctuations in repeaters and usually in switches can generally be reduced to acceptable levels by the use of statistical techniques amenable to low cost devices. Modern switches implement level 2 protocols that can further reduce fluctuations based on the priority of Sync messages [3]. It is also possible to design switches incorporating IEEE-1588 boundary clocks that logically remove them from the IEEE-1588 communication paths. This technique effectively produces direct synchronization between the clocks in the end device and the switch.

Further details on these mechanisms and other features of IEEE-1588 beyond the scope of this paper may be found in the standard itself or on the IEEE-1588 web site [5].

## **ORIGINAL CLOCK OSCILLATORS**

The clock oscillators used in the nodes of the original design are low cost devices available from numerous vendors. Known vendors and their part numbers are:

American KSS Inc. (Kinseki), Fremont, CA, USA	FXO-31FL 40.00MHZ
CTS Corp. 171 Covington Drive, Bloomingdale, IL 60108	CB3LV-3C-40.0000-T
Daishinku/KDS America Corp. Marietta, GA, USA	DSO751SV-40.000MHZ
Fox Electronics, Fort Myers, FL, USA	F4100 40.000MHZ
Kyushu Dentsu, Omura, JP	MIN30A-T 40.000MHZ

The specifications on the above units are all substantially identical. Salient specifications are:

Frequency (nominal)	40 MHz
Freq. Stability	+/-100 ppM
Supply Voltage	3.3+/-0.3 Vdc
Output Level	CMOS
Operating Temp.	-10 to +70 deg. C
Package	Ceramic Surface Mount 1.5 x 5.0 x 7.5 mm Max

## HIGH STABILITY OSCILLATORS

The oscillator that was substituted for the original low cost, modest performance unit is the Agilent Technologies 10811D, an oven controlled, high performance, highly stable device. This unit has been discontinued as a catalog product, but comparable units are available from other manufacturers.

## INTERFACE FOR HIGH STABILITY OSCILLATORS

Since the clock signal required by the circuitry in the node is a 40 MHz logic level signal from 3.3 V CMOS, and the output of the high stability oscillator is 10 MHz ac coupled sine wave, a frequency multiplication must be done to interface the two. The means chosen for this function was a phase locked loop with a frequency divider in the feedback path. The Philips NE564 chip was chosen to provide the phase detector and 40 MHz voltage controlled oscillator functions, while a 74LCX74 dual D type flipflop performs the  $\div 4$  frequency division. Additionally, a 74LCX04 hex inverter is used for buffering and output driving.

The bandwidth chosen for the NE564 loop was based on rough estimates of the noise of the phase detector and VCO versus the 10811D phase noise, and the need to smooth the ripple from the phase detector output. The noise considerations alone suggest a loop bandwidth of  $\sim 10$  MHz to maximally suppress VCO noise out to the crossover frequency.

However the phase comparator, a Gilbert cell double balanced multiplier, will have a strong 20 MHz ripple component in its output, and any useful filtering of this ripple will incur undesirable phase shift at the loop crossover frequency.

Hence a compromise choice was made to set the loop bandwidth at 1 MHz, for the  $\times 4$  interface PLL.

Breadboards of these interfacing functions were constructed and attached to two of the 10811D oscillators. From the interface breadboard, a 24 inch coax cable carries the 40 MHz signal to the node where the original surface mount low cost oscillator had been replaced with a coax connector. A photograph of this interface (mounted on an oscillator) is shown in Figure 2, and a schematic diagram is shown in Figure 3.

## OSCILLATOR MEASUREMENTS

Since the 40 MHz signal on the breadboards is divided down to 10 MHz for phase locking, the divided-down 10 MHz signal carries any instabilities that were incurred in the phase locked loop, and is buffered out for testing.

A third breadboard was constructed with one of the low cost oscillators followed by the  $\div 4$  frequency divider and the buffering IC, identical in that area to the other breadboards. The resulting 10 MHz signal can then be measured in a 10 MHz stability test system and compared with the 10 MHz performance of the 10811D oscillators and with the 10 MHz signal derived from dividing the 40 MHz signal that is locked to the 10811D output. This last signal will carry all the instabilities that the 40 MHz signal has.

Measurements of  $\sigma_y$ , defined as the two-sample deviation of fractional frequency fluctuation, at 10 MHz can be expected to yield the same result as if the measurements were made directly on the 40 MHz signal because the variable being measured is normalized versus the carrier frequency. It is true that any instability introduced by the divide-by-four circuitry will be included in the measurement at 10 MHz. Compared to the other instabilities of concern in this project, the contribution of the divider circuit can be considered negligible.

Tests of oscillator stability were conducted at Agilent's Santa Clara site using the production test systems that are normally used for testing quartz crystal oscillators and atomic frequency standards. These tests covered three units:

1. CTS CB3LV-3C-40.0000-T (surface mount) with  $\div 4$  divider
2. Agilent 10811D SerNr 704409 with X4 PLL and  $\div 4$  divider
3. Agilent 10811D SerNr 800764 with X4 PLL and  $\div 4$  divider

Time domain stability as Allan Deviation ( $\sigma_y(\tau)$ ) was measured on all three from  $\tau = 0.1$  second to 2000 seconds on the 10 MHz output of the  $\div 4$  divider. These are plotted together for comparison in Figure 4.

Phase noise as  $\mathcal{L}(f_m)$  was measured on the two 10811D units from  $f_m = 1.0$  Hz to 100 kHz on both the 10811D oscillator directly and on the 10 MHz output of the  $\div 4$  divider, Figs. 5, 6, 7, 8. The phase noise of the low cost oscillator was not measured because its 32 Hz offset at 10 MHz (well within spec.) was too great to be accommodated by the phase noise test system.  $\mathcal{L}(f_m)$  is defined as the single sideband-to-carrier ratio of phase noise spectral density, with units of dBc/Hz.

From the plots of  $\sigma_y$  we can calculate the jitter of the one second pulses due to the oscillator instability alone. This is done by reading the value for  $\sigma_y$  at a chosen  $\tau$ , then multiplying as:

$$\sigma_x = \tau \times \sigma_y$$

The averaging time,  $\tau$ , has the significance of being the elapsed time between successive measurements of time error, or the response time of a tracking loop.

Averaging Time $\tau$ , seconds	Jitter, $\sigma_x$ , nanoseconds	
	Low Cost	10811D
1	4	0.006
10	90	0.040
100	1400	0.700

From these results it is easy to see that for a tracking loop having a response time of a few seconds, the low cost oscillator can contribute several tens of nanoseconds of jitter, while the 10811D will remain well below one nanosecond.

## TIME SYNCHRONIZATION SERVO LOOP

The hardware that supports the IEEE 1588 implementation is an integrated circuit that performs the timekeeping function and communicates through an Ethernet port with other like nodes. The standard specifies how one of the node clocks is designated a master, and all the others are slaves. The slave clocks communicate with the master clock and synchronize themselves to the master. This is done by measuring and accounting for the propagation delay between the two clocks, then servoing the slave clock to agree with the delay-adjusted master. This servo loop is in fact an entirely digital phase-locked loop consisting of a time error measurement that is analogous to a phase detector, and a function that is the equivalent of a voltage controlled oscillator (VCO). There are also loop control functions to process the error signal, forming an integral and applying it and a proportional version to the VCO to close the loop.

## VCO RESOLUTION

The “VCO” function in the node is provided by a programmable adder that adds 24, 25 or 26 nanoseconds at each occurrence of the 40 MHz clock. This yields, on average, a phase step of +1, 0 or -1 nanosecond. A register controls the number and rate of these

phase steps, giving a resolution, averaged over hundreds of seconds, of  $1.49 \times 10^{-10}$ . Since this is actually a numerically controlled oscillator, we will refer to it as an NCO.

Note however that adding an extra count to the nanosecond counter has no effect on the phase of the output until 24 more adds have occurred, at which point the next one-second pulse will occur 25 nanoseconds earlier than it otherwise would have without the adds. Another way of saying this is that the granularity of phase adjustment is 25 nanoseconds.

The resulting granularity of frequency arises from this quantization step size of  $(25 \text{ ns}) \div (1 \text{ s}) = 25 \times 10^{-9}$ . This characterizes the way the NCO operates over a time span of a few seconds.

### **PHASE DETECTOR RESOLUTION**

The “phase detector” in this loop is the time error counter which is active every 2 seconds, and also has a quantization of 25 ns. Therefore the time error must reach 25 nanoseconds before the error counter reports anything other than zero time error.

### **CONTROL LOOP PARAMETERS**

With small or no frequency offset of a node clock (relative to the master clock), the best that such a time tracking servo can do is to hunt slowly, between +25 and -25 ns, plus some additional error accumulated between the time that the actual error reaches 25 ns and the time that the error reading is taken, on average one second; plus further additional error accumulated between the time that the error reading is taken and the time that the correction takes effect, a minimum of one second.

The original settings for the loop parameters were:

Proportional,  $P = 2$

Integral,  $I = 1/2$

With these settings, the loop gain was high, the bandwidth wide, and the response time fast, a desirable condition during startup to speed acquisition of lock and drive large initial errors to small values in a time span of less than a minute.

However with the loop so responsive, the perturbations caused by the granularity of the measurements and the delay between computing a change and having the result of the change appear in a measurement, caused the loop to hunt over a wide enough span to make the resulting standard deviation (from loop hunting alone) to be around 30 ns.

This finding suggested a significant reduction in loop responsiveness, even at the expense of startup time. The new loop parameters chosen were:

Proportional,  $P = 1/2$

Integral,  $I = 1/16$

Subjectively the startup dynamics seemed to take less than twice as long as with the earlier settings, and the hunting behavior was greatly improved. The standard deviation measured 13 ns, fully explained by the fact that the granularity of both the error measurement and the time phase adjustment occur in steps of 25 ns.

The amount of improvement of the hunting behavior is well summarized by the standard deviation, even though the hunting pattern is not at all random, being highly deterministic.

For a pair of nodes connected directly to each other and isolated from any other network connection, with both nodes having 10811D oscillators, the standard deviations were:

Servo Parameters		Clock Error Standard Deviation, $\sigma_x$ , nanoseconds
Proportional	Integral	
2	1/2	39
1/2	1/8	15
1/2	1/16	13

The first condition was the original setting of the loop and the last condition was the next thing tested. With the integral = 1/16, the settling time constant after startup was about 40 minutes. This result led to trying the middle of the settings with the integral = 1/8, where the settling time constant was about 20 minutes.

## NODE MEASUREMENTS

The setup for making measurements of the accuracy of time synchronization consists of the two nodes under test, an oscilloscope, an Ethernet switch, hub or direct cable, an Agilent 5372A Frequency and Time Interval Analyzer with Agilent 1145A Dual Active Probes. The 10811D oscillators and the interface boards are powered by a group of Agilent E3610A Power Supplies. Figure 9 shows a photo of a prototype node and gives a block diagram of the measurement setup. Figure 10 is a photo of the setup.

Although a boundary clock is not included in any of the measurements reported here, a photograph of a prototype boundary clock is shown as Figure 11. A boundary clock is necessary to coordinate synchronization between subnets that are separated by a router or other device that does not pass the IEEE 1588 messages.

The measurements are made by comparing the time of occurrence of the one pulse per second outputs of the nodes. These pulses are monitored on the oscilloscope to verify normal behavior, and are measured as to the time interval between the occurrences of the pulse from one node and the pulse from the other node under test. This gives a direct measure of the lack of synchronization between the two nodes. The data file from a series of measurements is downloaded from the time interval analyzer to a desktop computer via GPIB, and the data reduction computations and plotting performed by Excel. The time interval analyzer has as one of its features the ability to measure pairs of such pulses consistently as the pulses change from a state of "A leads B" to a state of "B leads A" and vice versa, with the appropriate sign change in the measured time difference. The first tests of the node clocks with 10811D oscillators were conducted with two identical units

connected only to each other via a cross-over Ethernet cable. An identical test of nodes having low cost oscillators was run for comparison.

The test of the unmodified nodes (with low cost oscillators) showed a random-appearing wander of delta-t, reaching peaks of 100 ns many times per hour, and a corresponding sigma of 34 ns.

The test of the nodes run from the 10811D clocks showed a much more deterministic, substantially periodic patterns of delta-t error. The shape of the excursions was quite similar over and over, being a ramp of about 50 ns over about 16 seconds of elapsed time. This corresponds to a frequency offset between the two clocks of  $3.1 \times 10^{-9}$ .

Tests were also run with the nodes connected through a hub and through a switch, with both original and modified servo settings. Both the hub and the switch were isolated from any other network to avoid unknown and unpredictable perturbations. The only messages through the hub and switch subnets were the synchronization messages originated by the nodes. The hub is an HP ProCurve 10Base-T Hub 8, model number J4090A, and the switch is one of five HP ProCurve Switch 10/100Base-T modules, model number J4111A, mounted in an HP ProCurve Switch 4000M frame with a model number J4121A Switch Engine Module.

Some earlier versions of the node firmware included a measurement data editing algorithm that would reject and ignore error readings that were unexpectedly large, “outliers.” This outlier rejection was disabled for all the tests reported here.

The raw data from the measurements is plotted in Figures 12 through 23. The test conditions for this series of figures are:

Figure	Oscillators	Tracking Loop	LAN Connection
12	Inexpensive	Fast	Direct
13	Inexpensive	Fast	Hub
14	Inexpensive	Fast	Switch
15	Inexpensive	Slow	Direct
16	Inexpensive	Slow	Hub
17	Inexpensive	Slow	Switch
18	10811D	Fast	Direct
19	10811D	Fast	Hub
20	10811D	Fast	Switch
21	10811D	Slow	Direct
22	10811D	Slow	Hub
23	10811D	Slow	Switch

Many of the plots have a smoothed trace plotted as a lighter color; this shows a running mean with a four minute averaging time. Several of the plots show large swings at the start of the test run. On these the data logging was started before the synchronization was enabled in order to see the acquisition and settling behavior in its entirety. The settling periods were excluded from the computations of standard deviation and mean.

## STANDARD DEVIATION OF ERRORS

The more or less random scatter of the time synchronization errors can be characterized by computing their standard deviations. This data is shown in the following table:

Conditions: No rejection of outliers  
Hub and switch isolated from net

	Time Synchronization Error, $\sigma_x$ , nanoseconds	
Original Servo Parameters P = 2, I = 1/2	Low Cost Oscillator	10811D Oscillator
Direct	34	39
Hub	75	46
Switch	140	133
Best New Servo Parameters P = 1/2, I = 1/8	Low Cost Oscillator	10811D Oscillator
Direct	76	15
Hub	80	42
Switch	123	92

As the table shows, the more stable oscillator resulted in better time synchronization performance in all cases except the direct connection with the original servo parameters, where the instability of the low cost oscillator seemed to assist in a reduction of the hunting behavior of the faster loop, much as the addition of dither is sometimes used to aid a coarsely digitized system.

## MEAN ERRORS

The total time synchronization error consists of not only the more or less random deviations, but also any non-zero mean values. The mean values corresponding to the above deviations are shown in the following table:

Conditions: No rejection of outliers  
Hub and switch isolated from net

	Time Synchronization Error, mean, nanoseconds	
Original Servo Parameters P = 2, I = 1/2	Low Cost Oscillator	10811D Oscillator
Direct	-28	-1
Hub	-27	-7
Switch	-49	-14
Best New Servo Parameters P = 1/2, I = 1/8	Low Cost Oscillator	10811D Oscillator
Direct	-21	+5
Hub	-32	+10
Switch	-21	+5

No conclusions are drawn as to the causes of these offsets. Observations are that the change in servo parameters and the various network paths seem to have little effect, and that the more stable oscillator seems to have a very beneficial effect in reducing this error term. The reasons for this are not at all clear.

## DO THE VARIANCES ADD?

Comparison of the standard deviation among the various cases leads to the question of whether each element of a particular test setup contributes a predictable amount of instability. If the statistical processes were truly random, normally distributed, and independent, then we could expect that the variances of each would add to form the variance of the overall system. Since *none* of these three conditions is satisfied, we should not expect addition of variances to be an accurate model for these results. However, applying the hypothesis of additive variances does lead to some useful reasonable approximations to the observed effects. The central limit theorem, sometimes called the law of large numbers, may apply here.

To apply this hypothesis of additive variances, first take the cases of low cost oscillator, slow servo loop, and compare the standard deviation of the direct case with the case including a switch:

Direct 76 ns

Switch 123 ns

$$(76)^2 + (\text{Switch alone})^2 = (123)^2$$

Switch alone = 97 ns

Working this same calculation with the 10811D oscillator:

Direct 15 ns

Switch 92 ns

$$(15)^2 + (\text{Switch alone})^2 = (92)^2$$

Switch alone = 91 ns

Here the two results for the switch alone are close enough to guess that the switch alone might be expected to contribute around 94 ns of one-sigma uncertainty to any path.

A similar analysis of the cases involving the hub yields results that are not as consistent.

The data collected here also indicate that the individual elements contribute differing amounts of instability depending on the parameters of the tracking servo. For the fast servo loop parameters, the switch alone with the low cost and 10811D oscillators respectively, gives 136 and 132 ns. This suggests that the switch may be expected to contribute about 134 ns when working with the fast loop, versus the 94 ns with the slow loop.

An attempt was made to apply this analysis to the quantization noise and tracking behavior of the node alone, and the low cost oscillator alone, based on the assumption that the 10811D oscillator contribution is negligible. Unfortunately the method fails to yield meaningful results, probably because the data points are not normally distributed.

It appears that the instability of the hub alone is influenced by the oscillator performance as well as the choice of servo parameters. It is not at all clear why the hub performs better

with a fast loop with the 10811D oscillator but better with a slow loop with the low cost oscillator. The many interactions of oscillator instability with loop dynamics might explain this, with a better understanding of the details of those interactions.

The following table presents the results for comparison:

	Error, $\sigma_x$ , ns	
	Tracking Loop	
	Fast	Slow
Direct with 10811D Oscillator	39	15
Direct with Low Cost Oscillator	34	76
Hub alone, with Low Cost oscillators	67	25
Hub alone, with 10811D oscillators	24	39
Switch alone	134	94

Summarizing these results we can say that the fast tracking loop parameters seem helpful for the low cost oscillator, and are clearly sub-optimal for a more stable oscillator. The slower loop parameters allow the quantization noise to have less effect, as well as working better with the instabilities coming from the switch to give a 40% lower result. The two different loop settings had oscillator-dependent effects on instabilities originating in the hub. In particular with the more stable oscillator, the dynamics of the tracking are more of a hunting behavior, not a normally distributed random variable, and these analytical techniques cannot be expected to offer precise answers.

## **OPPORTUNITIES FOR FURTHER DEVELOPMENT**

This project explored the improvements enabled by a much more stable clock oscillator, and quantified the limitations imposed by other sources of instability as they were made clear by the elimination of the oscillator as a significant source of error. Since the time tracking servo parameters are all software-implemented, adjustment of these parameters was tested but only to a “first look” exploratory level. There may be opportunities for further tuning of the servo behavior to better deal with instabilities of network hardware (hubs, switches, routers) and to complement other changes in the node itself.

There is a range of hardware that should be considered in the next phase of testing; 100Base-t links, fiber optic links and others. The testing should begin to include realistic traffic loading on the links. Additional investigation into the use of various statistical methods for reducing the effects of outliers and biases is warranted.

Since the concept of a boundary clock to transfer time between subnets and around routers is defined in the standard, this is an important area that deserves investigation.

The most obvious avenue for improvement in the prototype nodes is a move to 100 MHz for the basic clock frequency, 2.5 times the present frequency. The advantages of such a move are more precise measurements of time error, more precise adjustments of time phase, and more processor cycles available for all tasks. The disadvantages are that a hardware redesign cycle is required, likely higher power consumption and emitted interference. At this time 100 MHz low-power logic is a fully mature technology.

Another opportunity area is the oscillator. As the results indicate it would not be necessary to use an ultimate performance (and correspondingly high cost) oscillator, but that an order of magnitude improvement in the oscillator could be a good choice considering cost versus benefit.

A further development worthy of consideration is the concept of making the oscillator electrically tunable, a VCXO (voltage controlled crystal oscillator) and developing the control voltage using a DAC (digital to analog converter) from a register in the processor that accumulates a digital frequency correction factor. This has the significant advantage of completely removing the considerable instability incurred by phase-stepping the clock to maintain synchronization. Phase stepping could still be used during coarse acquisition of synchronization to speed startup, while using the smooth phase shifts of frequency control to achieve a non-granular tracking behavior.

Finally, the improvement in mean synchronization error seen with the more stable oscillator is an unexplained benefit that might yield further improvements with better understanding.

## **CONCLUSIONS**

The objective of characterizing time synchronization error performance of IEEE 1588 nodes with low cost oscillators versus high stability oscillators was met. Test results showed mostly consistent improvement with the better oscillator, and significant further improvement by adjusting the tracking servo parameters. The combined effects lowered the standard deviations of time errors by factors in the neighborhood of two. The removal of significant contributions of instability from the oscillators allowed for a degree of characterization of the instabilities of other elements of the node and network hardware which turned out to be comparable to or worse than the low cost oscillator.

## **ACKNOWLEDGEMENTS**

John Eidson of Agilent Laboratories and Joe White of the U.S. Naval Research Laboratory did the origination and planning of this project. John provided management and direction of the project, as well as serving as the major technical reference resource. Bruce Hamilton of Agilent Laboratories provided a great deal of assistance in this project. Bruce ran many of the tests, modified software and firmware to adjust

parameters, and gave invaluable insights in the many discussions and analyses of data. This project was funded by the U.S. Naval Research Laboratory to provide the services of Mike Fischer.

## REFERENCES

[1]: ‘Synchronizing Measurement and Control Systems’, John C. Eidson, R. William Kneifel II, Anatoly Moldovansky, and Stan P. Woods. Sensors Magazine, November 2002.

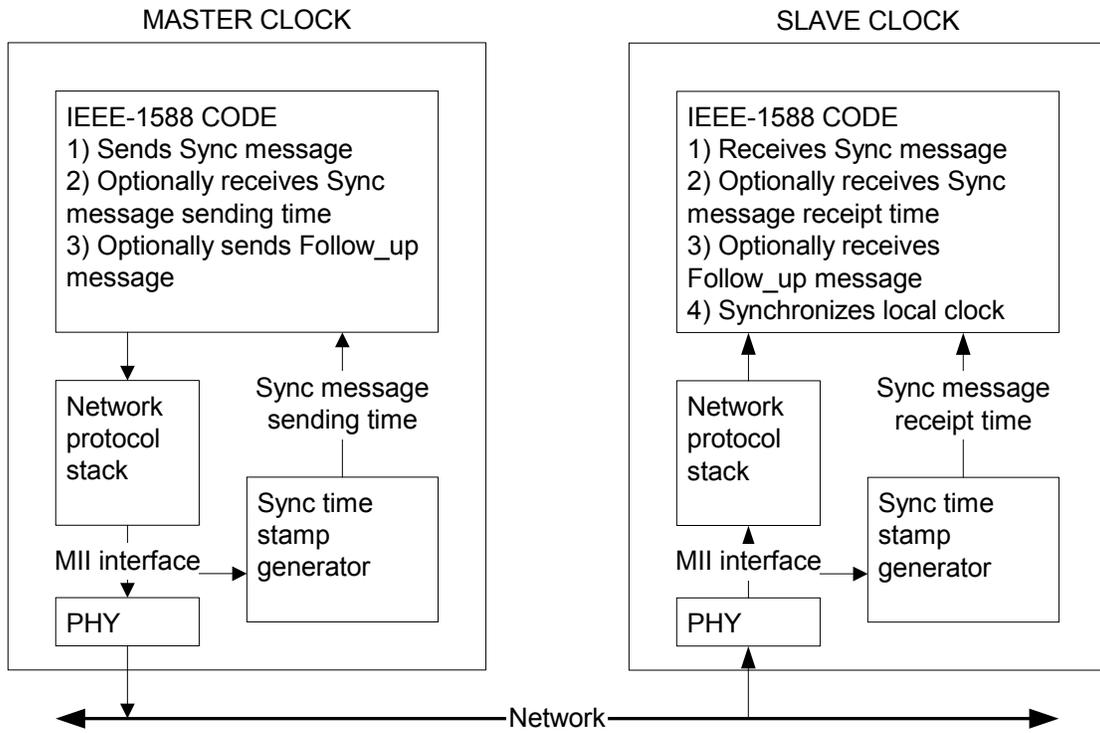
[2]; See <http://standards.ieee.org/>

[3]: ‘Utilization of Modern Switching Technology in EtherNet/IP™ Networks’, Anatoly Moldovansky, RTLIA 2002, 1<sup>st</sup> Int’l Workshop on Real Time LANS in the Internet Age, Technical University of Vienna, Austria, June 18, 2002

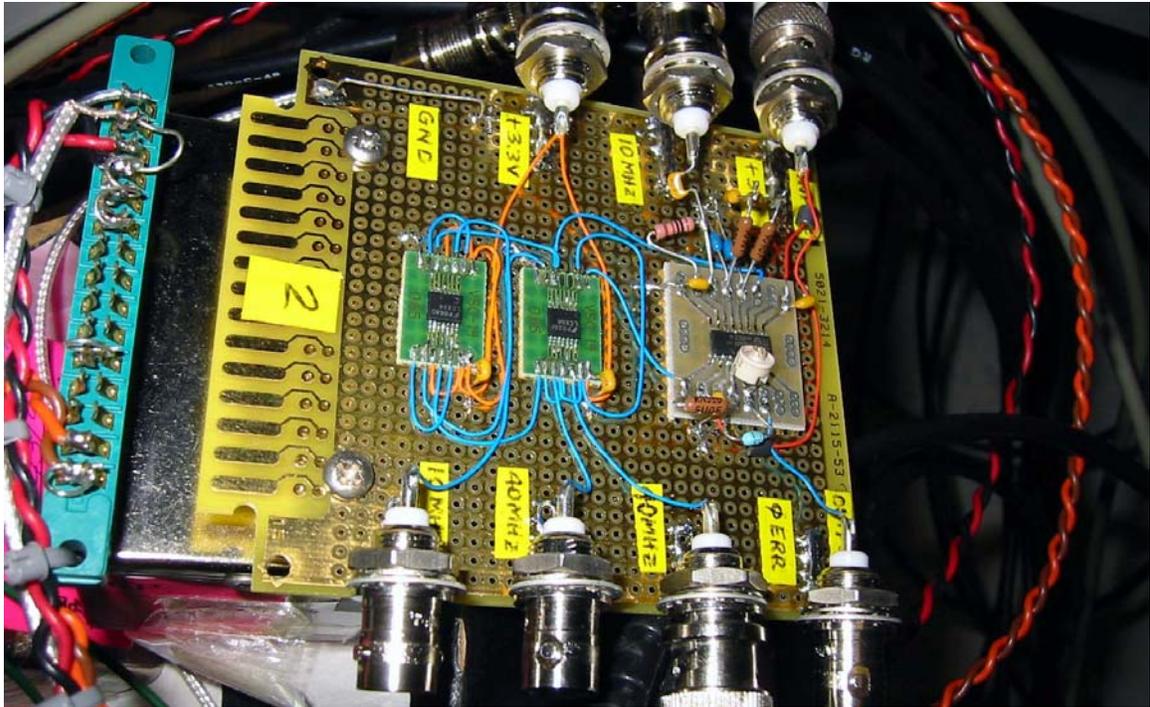
[4]: See <http://www.eecis.udel.edu/~ntp/> for details concerning the Network Time Protocol.

[5]: See ‘IEEE 1588 Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems’ at <http://ieee1588.nist.gov>

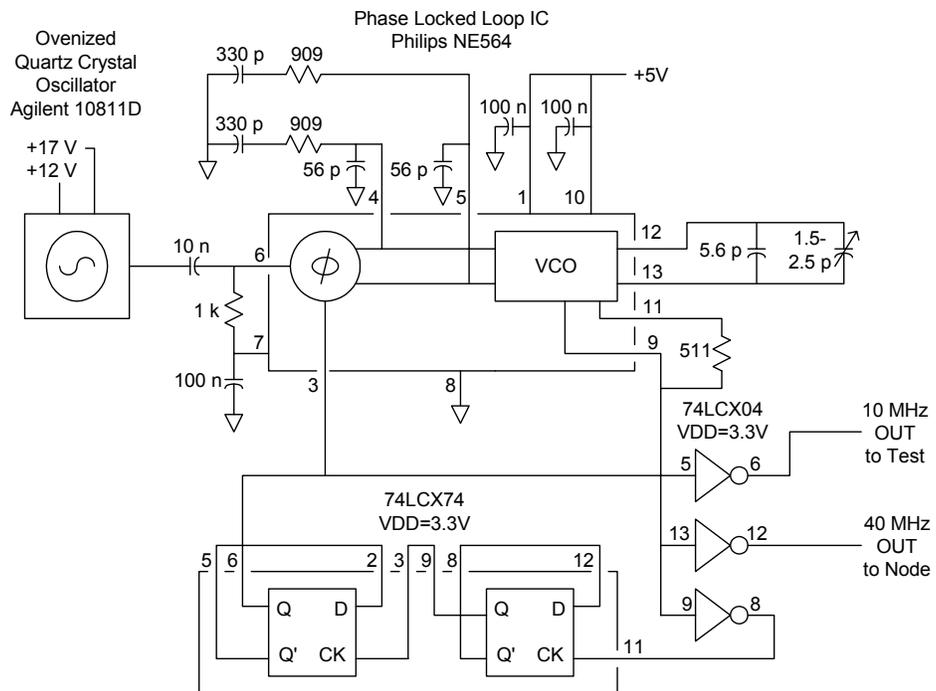
[6]: “Common Time Reference Technology for Systems Interoperability”, R. Beard and J. White, U.S. Naval Research Laboratory, Proceedings of the 32nd Precise Time and Time Interval Systems and Applications Meeting, 2000, pp 415-422.



**Figure 1: IEEE-1588 messages used in synchronization**



**Figure 2: High stability oscillator with interface PLL**



**Figure 3: Interface for high stability oscillators; PLL x4**

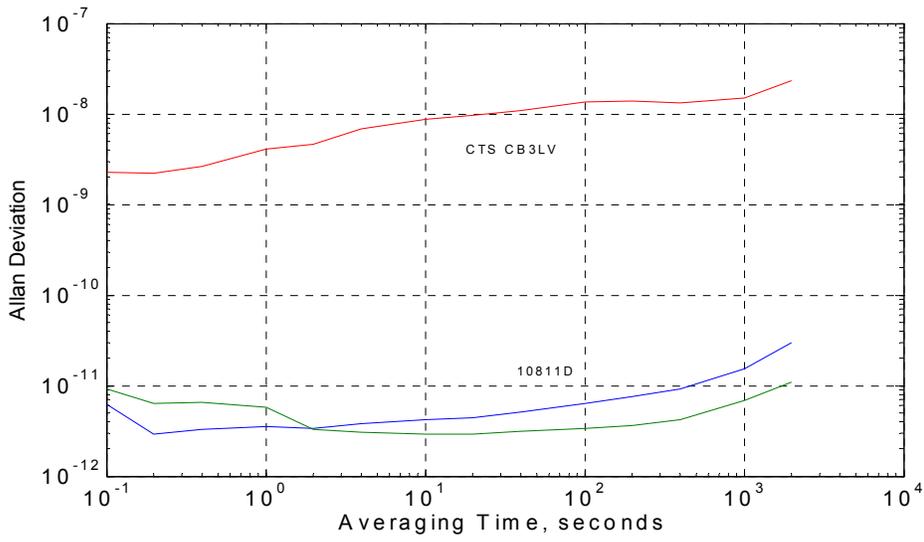


Figure 4: Time Domain Stability of the Oscillators

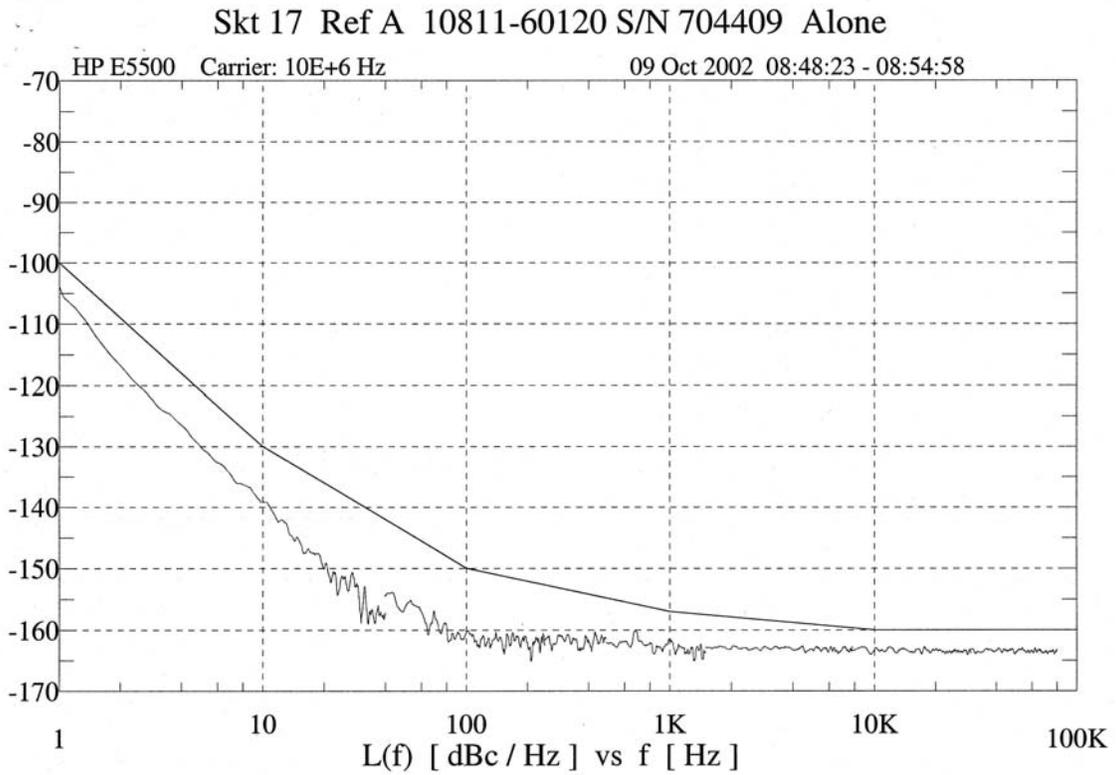


Figure 5: 10811D Oscillator 1, phase noise before  $\times 4$

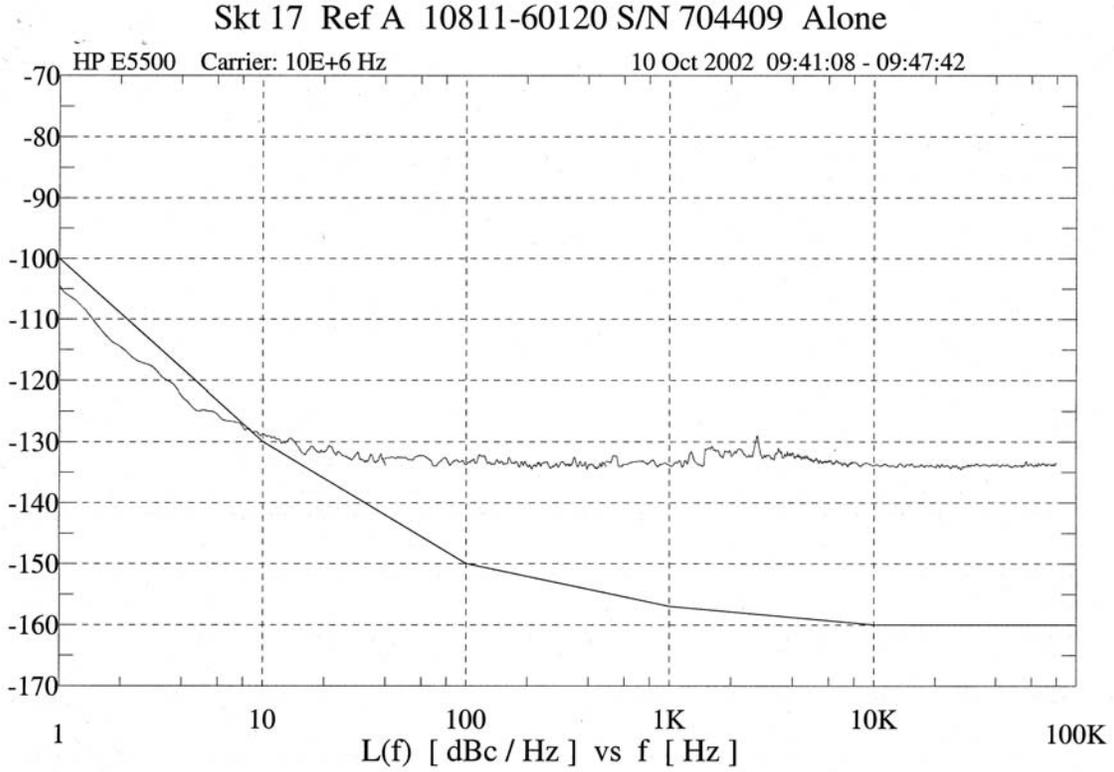


Figure 6: 10811D Oscillator 1 + interface, phase noise after  $\times 4, \div 4$

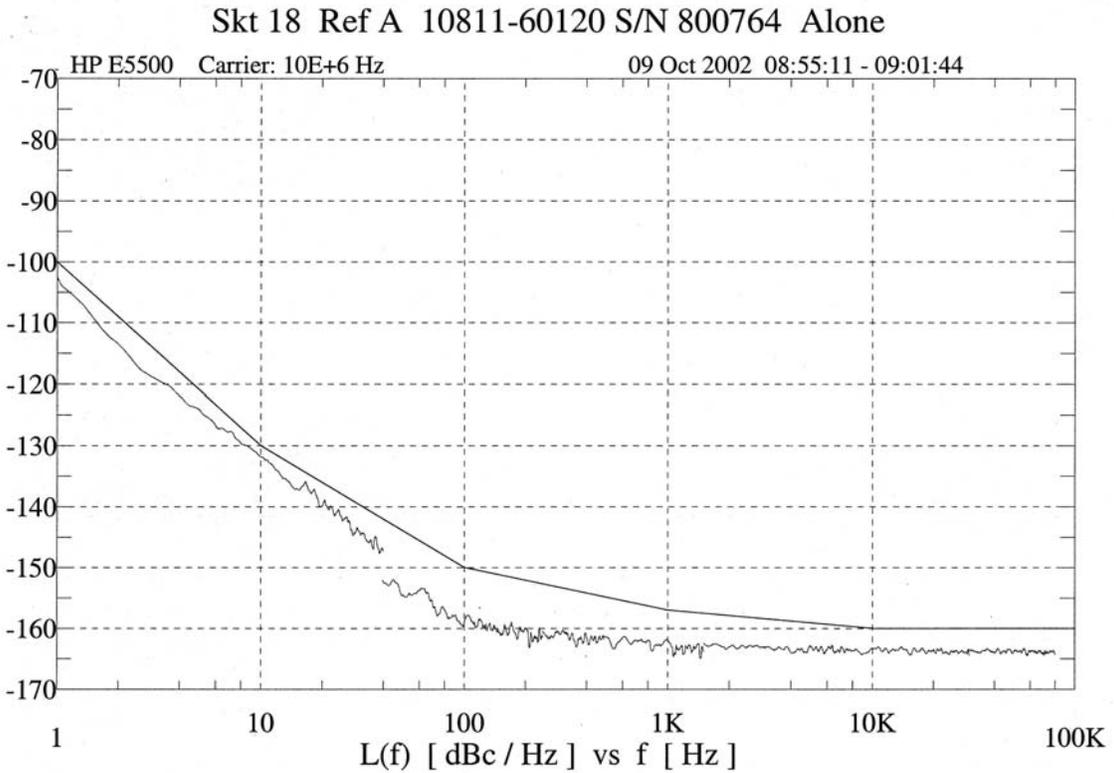


Figure 7: 10811D Oscillator 2, phase noise before  $\times 4$

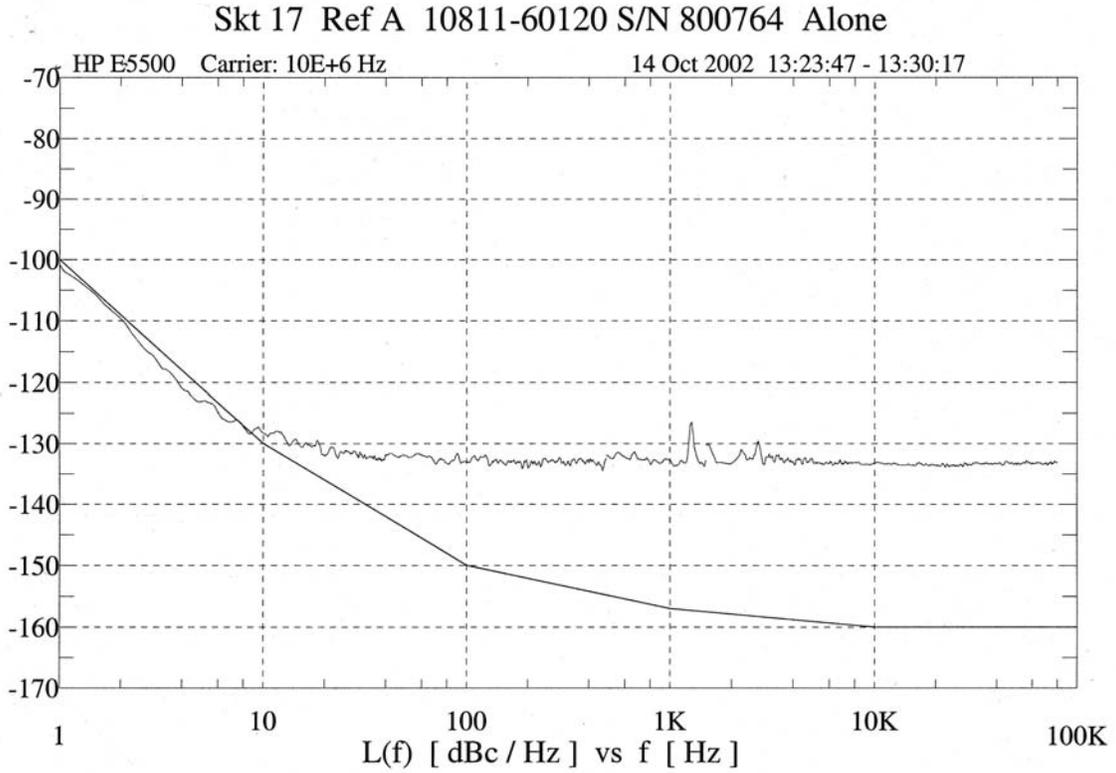


Figure 8: 10811D Oscillator 2 + interface, phase noise after  $\times 4, \div 4$

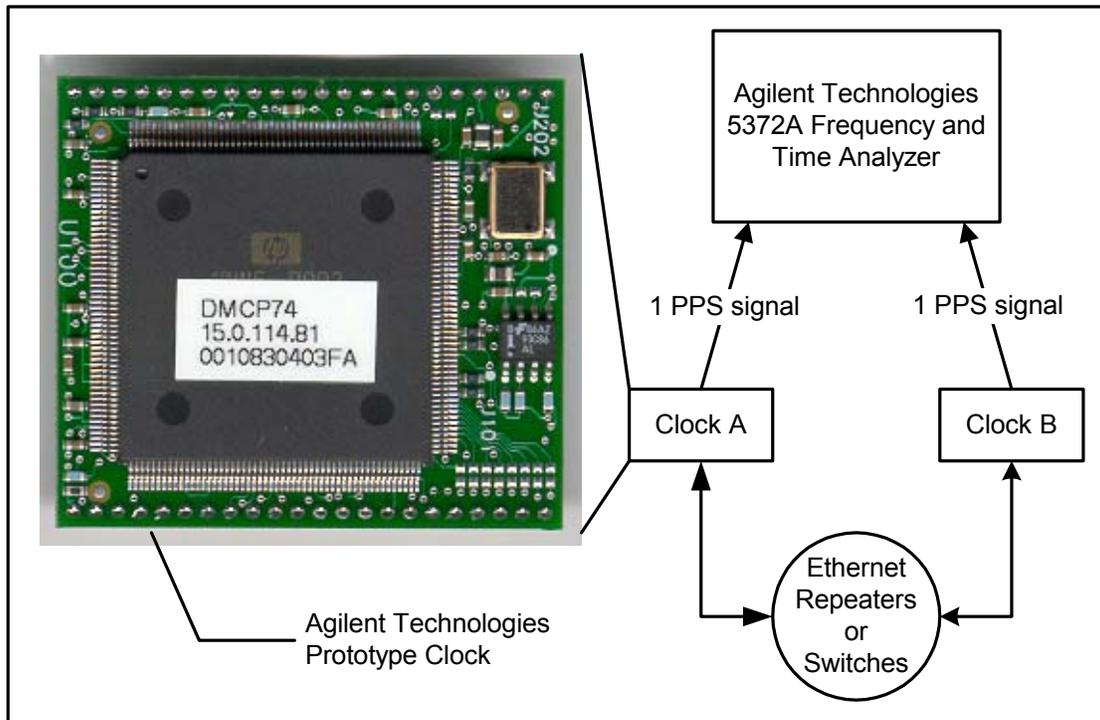
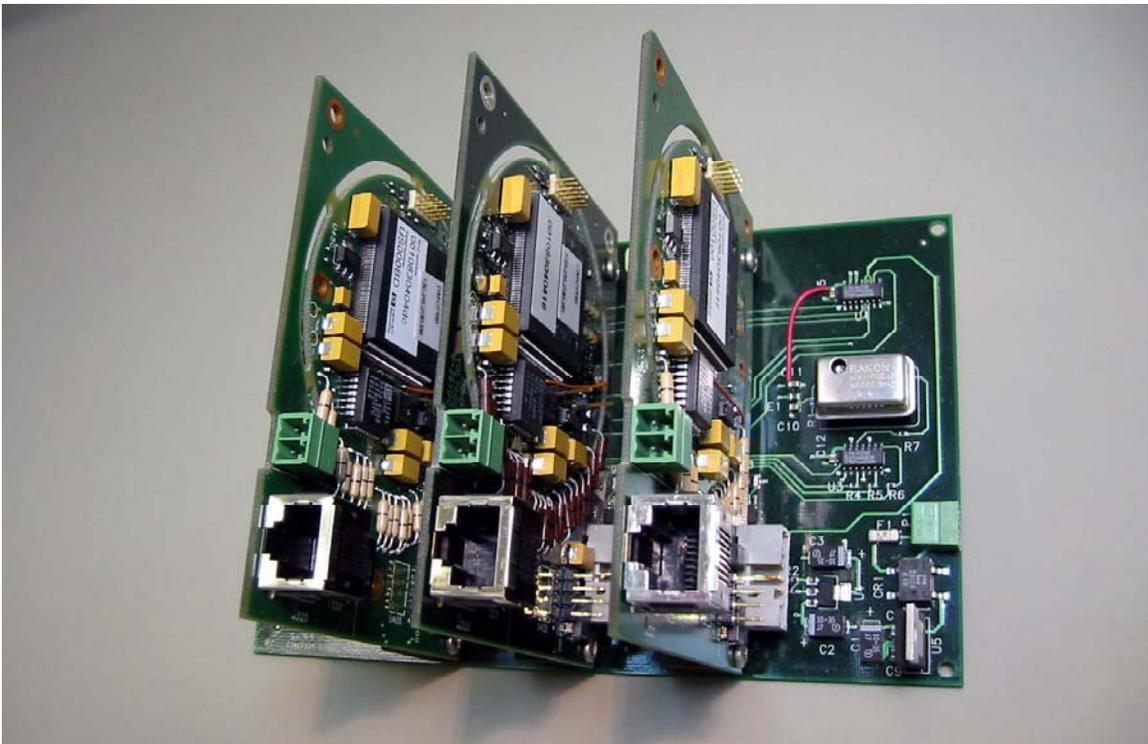


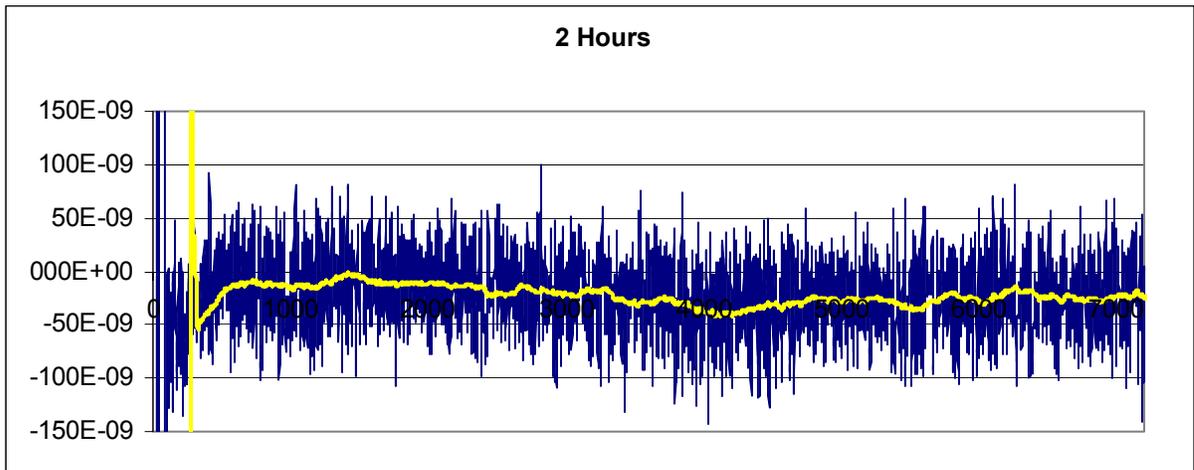
Figure 9: Measurement topology



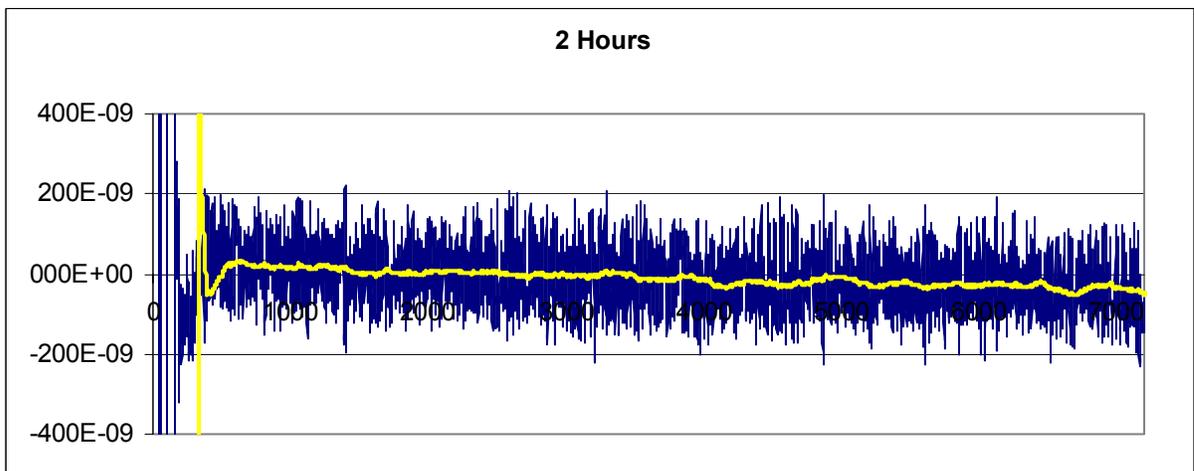
**Figure 10: Measurement setup**



**Figure 11: Prototype boundary clock with three Ethernet ports**



**Figure 12: Inexpensive oscillators, fast servo, direct connection**



**Figure 13: Inexpensive oscillators, fast servo, through hub**

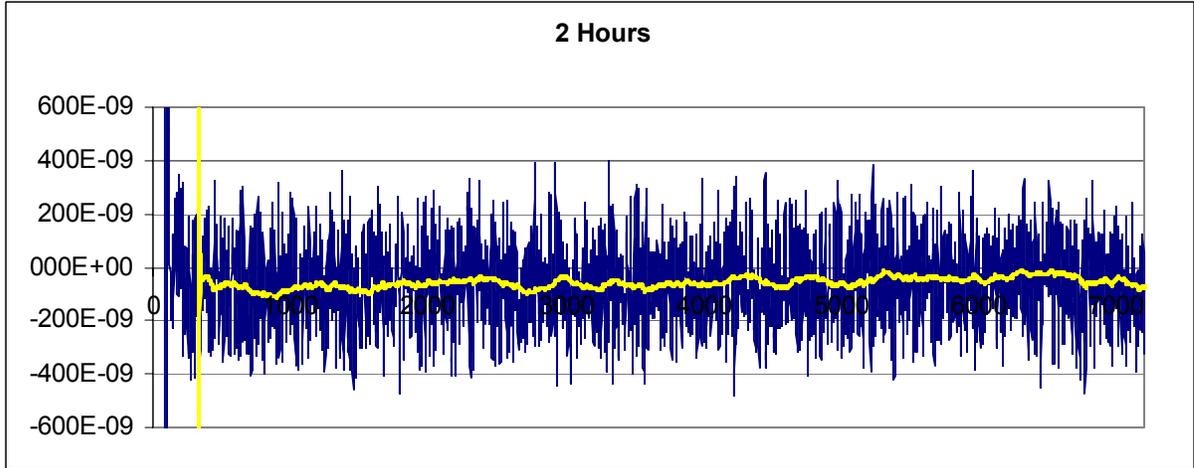


Figure 14: Inexpensive oscillators, fast servo, through switch

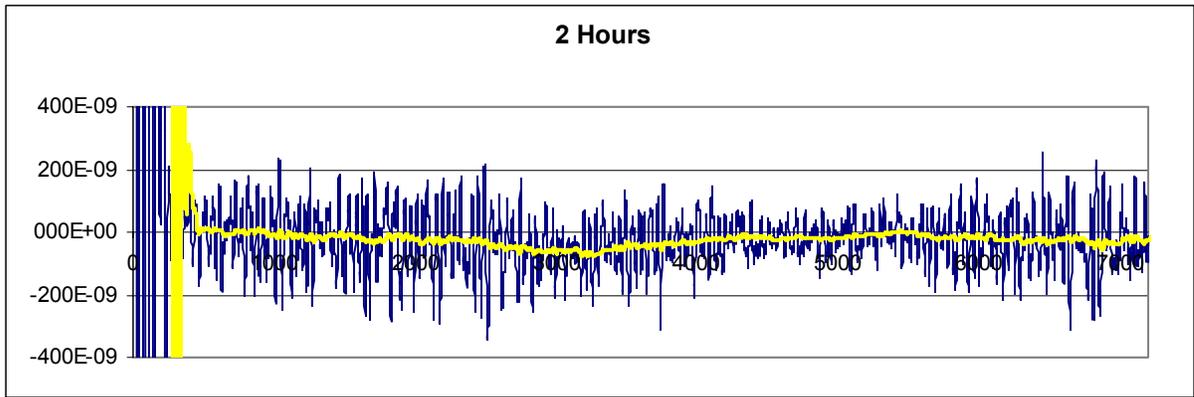


Figure 15: Inexpensive oscillators, slow servo, direct connection

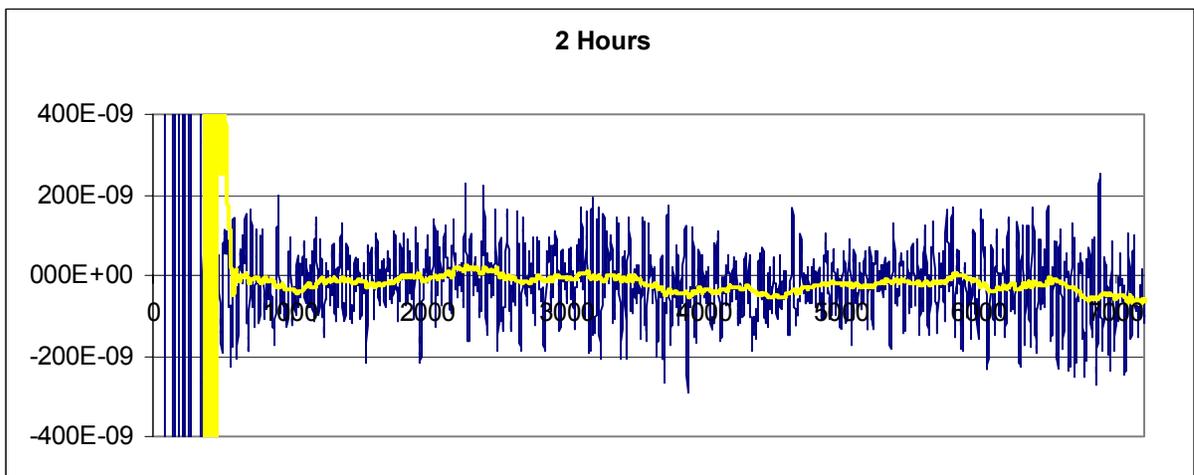


Figure 16: Inexpensive oscillators, slow servo, through hub

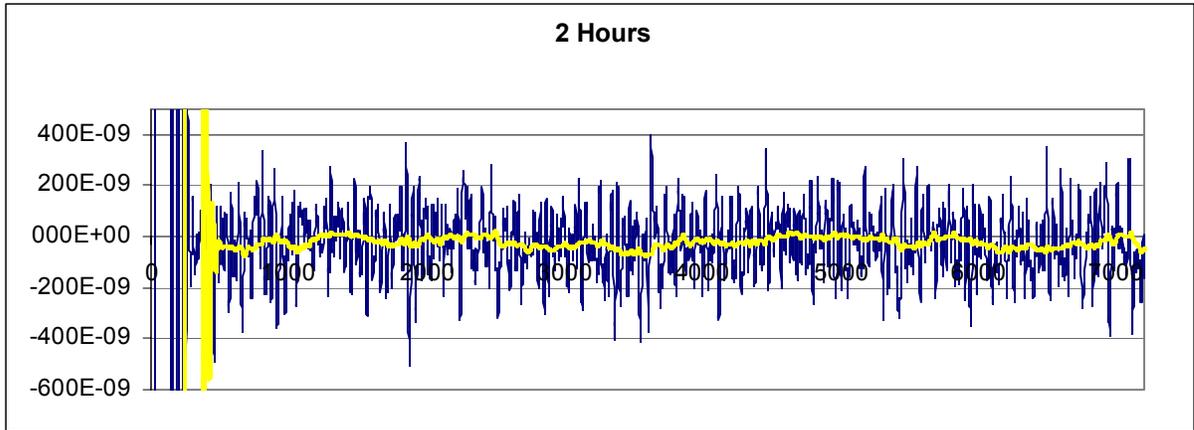


Figure 17: Inexpensive oscillators, slow servo, through switch

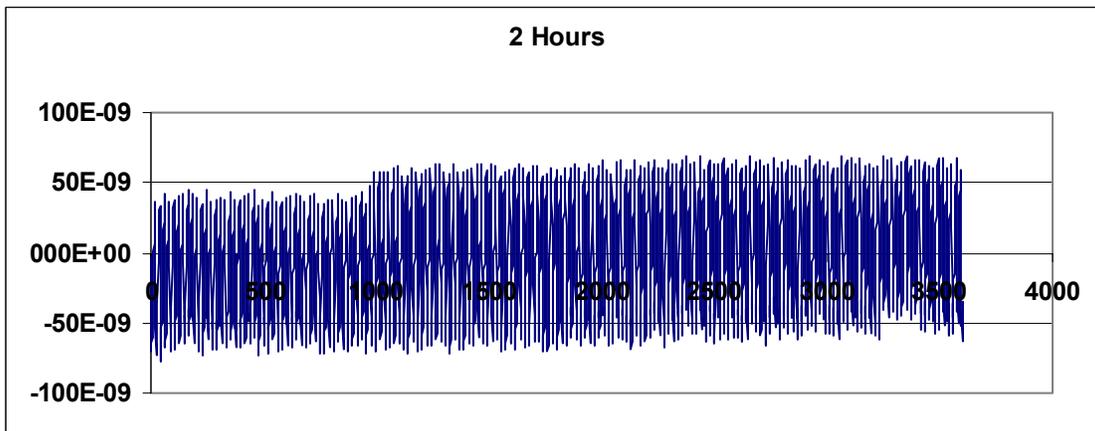


Figure 18: 10811D oscillators, fast servo, direct connection

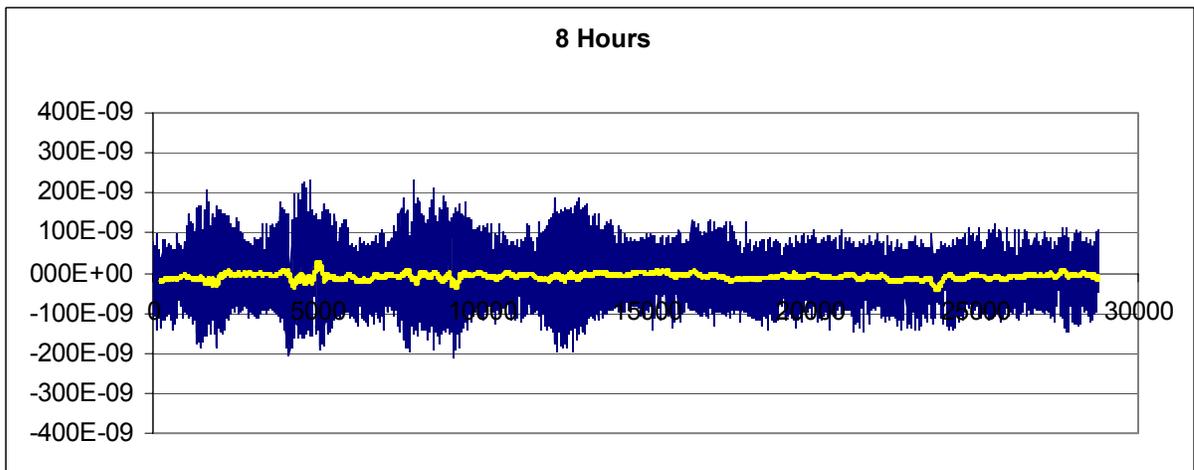


Figure 19: 10811D oscillators, fast servo, through hub

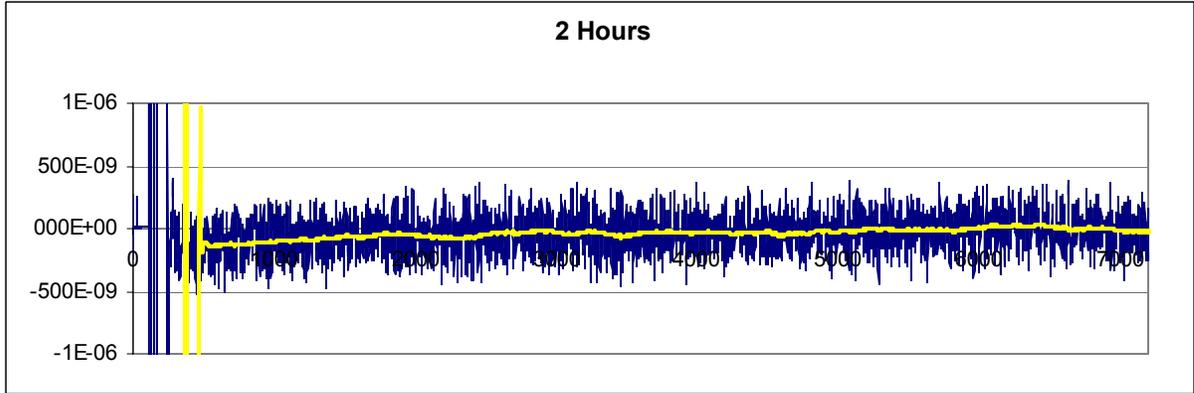


Figure 20: 10811D oscillators, fast servo, through switch

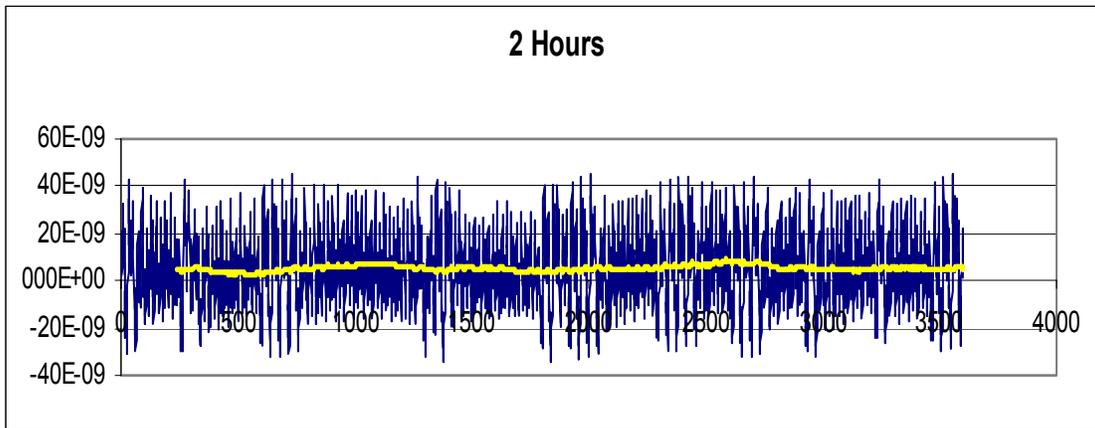


Figure 21: 10811D oscillators, slow servo, direct connection

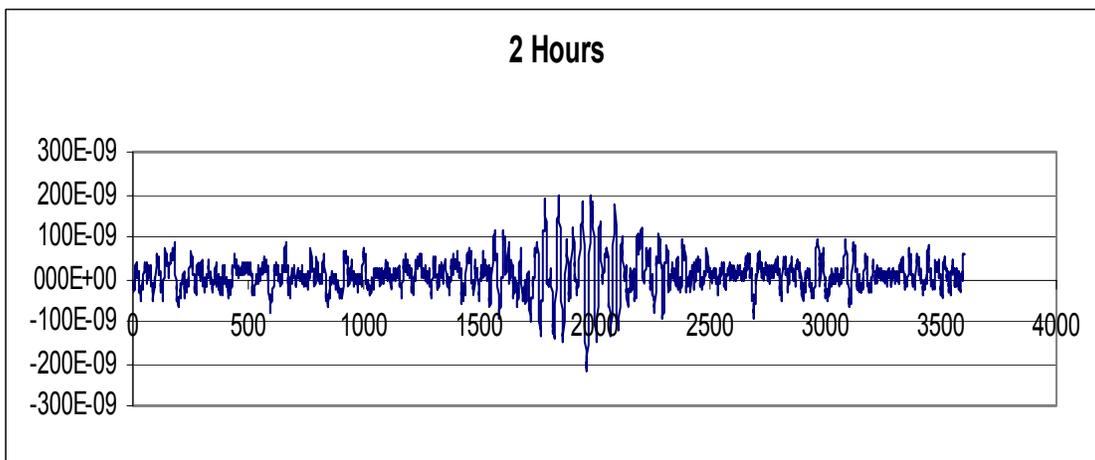


Figure 22: 10811D oscillators, slow servo, through hub

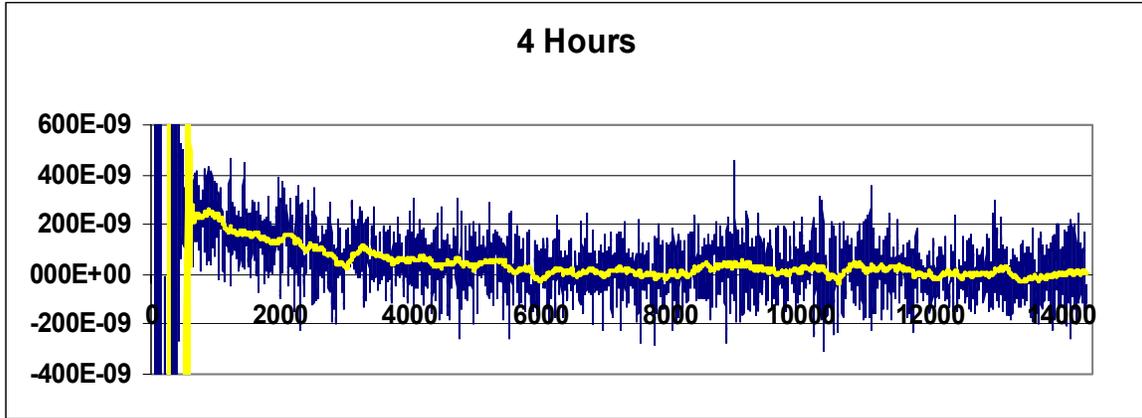


Figure 23: 10811D oscillators, slow servo, through switch