Oscillator Phase Noise:
Theory vs. Practicality

APP NOTE

Greenray

Modern electronics use time bases for a wide variety of applications. Communications systems rely on time bases for modulation and demodulation of data, GPS systems rely on them for accurate positioning, and a host of other applications rely on time bases to manage the flow of data within their system.

As the range of applications grows and the frequency increases, designers need time bases that are tighter in stability and lower in noise. As a result, oscillator designers need to continually push the limit of tight stability, low noise oscillator design.

One of the major issues facing oscillator designers is the phase noise phenomenon. Phase noise is an undesirable entity that is present in all real world oscillators and signal generators. It can cause distortion or complete loss of incoming information in traditional receivers, and high bit error rates in phase modulated applications. It is, therefore, necessary to understand and to quantify phase noise so that its effects on the higher level product are minimized.

Time is a concept all humans innately grasp as a function of life. It defines our world into days, months and years. When we look at a clock we know what time it is. However, do we really understand time and how it is defined? Like all measurements, time is measured with a degree of uncertainty. In engineering applications we use time bases that are imperfect by nature and add uncertainties to the applications that are driven by them. One byproduct of this uncertainty is phase noise. The inherent instability present in every oscillator manifests itself as a spectrum of noise around the frequency of the oscillator. This noise band is generally measured from the carrier to 1MHz away from the carrier. It is depicted as a graph of dBc/Hz vs. f(Hz), which shows how far down from the carrier power the noise power at a given frequency is.

In the following paper, many topics will be discussed including:

1. Different time bases and their properties
2. Crystal oscillator properties
3. Phase noise definition and calculations.
4. Ways to reduce phase noise in oscillator design
5. Real oscillator phase noise plots
6. Outside noise/interference effects on phase noise

The following information will show that phase noise is a manageable problem. Oscillator designers can work to minimize phase noise in oscillators, and system designers using the oscillators can better design their systems by choosing the correct oscillator. Outside noise and interference degrade the phase noise performance of oscillators, and engineers need to be aware of these factors so that their effects can be anticipated or avoided.

■ Time Bases

As engineers we are used to looking at our frequency counter and reading the display to determine what frequency we are operating at. But what is the accuracy of that measurement? The uncertainty of the measurement is dependant upon the accuracy of the time base the counter uses to drive its internal circuitry. This time base is not perfect and therefore skews the results of the frequency reading. For those applications requiring very precise frequency measurements the accuracy of this time base is of the utmost importance.

Time, like every other measurement man makes, has an uncertainty. However, time is the most accurate standard that mankind has produced. The second is defined as the resonant frequency of cesium 133 (Cs-133) which is 9,192,631,770 Hz [8].

Therefore, by measuring the vibrations of cesium, we have a standard for time. NIST (National Institute for Standards and Technology) is the keeper of the time standard for the United States. As of 2005 the NIST standard has an uncertainty of $5X10-16$ s, which means it would neither gain nor lose a second in 60 million years [9].

Cesium clocks, despite their accuracy, have several drawbacks that make them unsuitable for commercial electronics use. The cost of cesium standards is prohibitive for using them as a time base. Secondly, they are large. The NIST Standard fills a good portion of a room, and while smaller cesium standards are available they are not a portable item. Also, the cesium standard has a warm-up time. Keeping it on power is important, because a break in power could mean degradation in accuracy. Lastly, cesium standards use fuel - the cesium naturally depletes rendering the device effectively 'out of gas'.

Since cesium cannot be used effectively for commercial electronics or for most lab applications, other timing sources must be considered. Rubidium offers extremely accurate frequency counting similar to cesium, however it suffers from the same pitfalls. Being only slightly less accurate than cesium, the rubidium standards have a high cost associated with them. They are not portable devices and must be in a fixed location. The rubidium standards also have the same fuel depletion problem as cesium.

The choice for most consumer electronics is the quartz crystal-based oscillator. Quartz operates via the piezoelectric effect. Voltage applied to the crystal causes it to vibrate in a very stable and predictable way. By cutting and shaping the quartz a desired frequency of oscillation can be obtained.

Quartz has many advantages: it's cheap (compared to cesium and rubidium), it's small (crystals can be obtained in package sizes less than 3.2mm x 2.5mm), and it has a high Q (>500K for larger blanks). These traits have made quartz the choice for timing devices for many decades..

■ Quartz Crystal Oscillators

Crystal oscillators come in many types, shapes, and sizes. XOs (crystal oscillator) are a quartz crystal and driving circuitry. These bare bones clocks are cheap and small but offer limited accuracy as the crystal will wander in frequency approximately +/-30ppm over temperature.

TCXOs (Temperature Controlled Crystal Oscillators) use a compensation voltage to correct for the crystals natural temperature drift. This is accomplished with a classic thermistor resistor network or a polynomial generator. TCXOs are tighter stability over temperature (+/-0.25ppm can be obtained), and are small (2.0mmx2.5mm available) with low power consumption (2mA in some cases).

MCXOs (Microprocessor Controlled Crystal Oscillator) use a microprocessor to correct for the crystals natural temperature drift, by sensing the temperature of operation and using that data to correct the frequency of the oscillator. These oscillators can achieve stabilities of +/-0.1ppm over temperature, but have a slightly larger footprint, consume more power, and have degraded noise characteristics due to the microprocessor running in the oscillator.

CXOs (Oven Controlled Crystal Oscillators) and DOCXOs (double oven controlled crystal oscillators) offer the greatest stability that quartz crystal oscillators have to offer. By heating the circuitry inside, the crystal stays at an almost constant temperature, nearly eliminating the natural temperature drift of the crystal. Stabilities of parts in 10^{-10} over temperature are achievable, but with the tradeoff of package footprint (they measure at least 1 inch square), and power consumption (possibly over 1A).

Despite having a high Q, the quartz crystal oscillator is not perfect. Ideally a sinusoidal oscillator would produce a voltage as in Equation 1:

$$V(t) = V_O \cos(2 \pi f_o t) \quad (1)$$

Where V_O is the amplitude, f_o is the frequency, and t is time.

However, real world oscillators have some amplitude fluctuations and phase fluctuations present within them and behave as in Equation 2 [4].

$$V(t) = [V_O + \lambda(t)]\cos(2\pi f_o t + \phi(t)) \quad (2)$$

Where V_O is the Amplitude, f_o is the frequency, t is time, $\lambda(t)$ is amplitude fluctuation, and $\phi(t)$ is the phase fluctuation.

The frequency of the oscillator is affected by several factors: temperature, long-term drift, and short term instability. Quartz is quite sensitive to temperature, and the frequency generally drifts between +/-30ppm over the temperature range of the crystal. Long-term drift, also called aging, is a natural phenomenon of the quartz and is well understood. The aging characteristic is defined as:

$$f(t) = A(\ln(Bt + 1) + f_o \quad (3)$$

Where t is time in days and A,B,and f_o are constants determined from least squares fit (as per MIL-PRF-55310D) [7].

Figure 1

Typical TCXO Aging

3
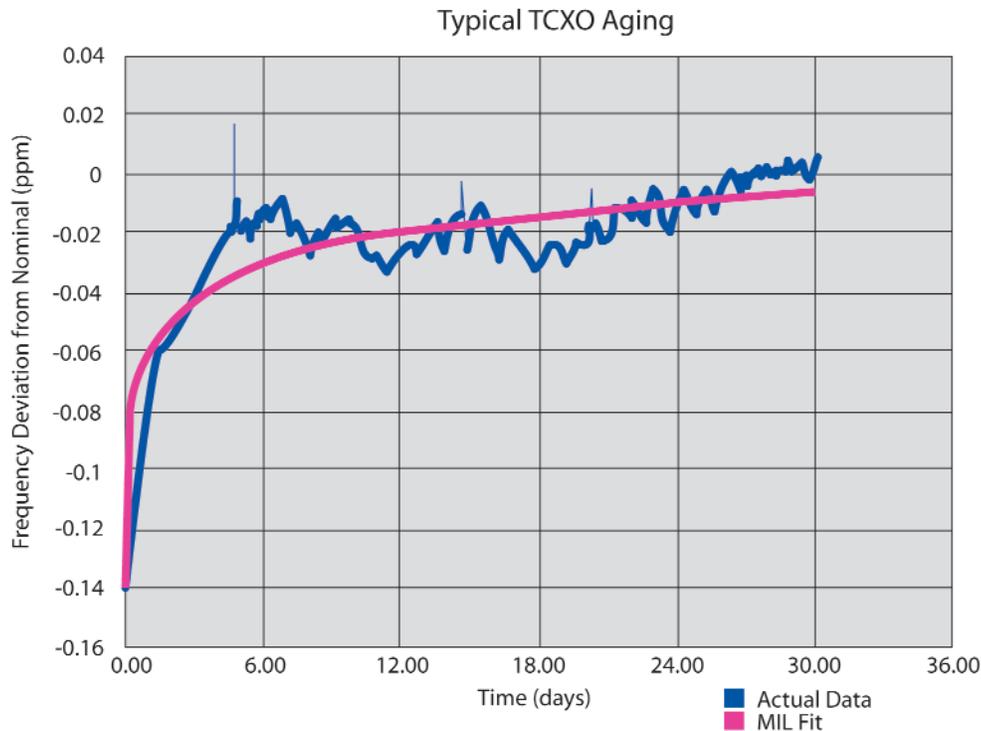
As cam be seen in Figure 1, the aging characteristic of quartz is a function that slows over time. This means that the frequency drift of the oscillator will diminish as time passes. This is a desirable effect in terms of long term performance. The oscillator will drift but that rate of change will slow and the oscillator in effect will become more stable.

Short term stability or Allan variance (AVAR) is a measurement of the short term frequency variations from the oscillator. Generally AVAR is specified relative to a particular gate time. For example, a 20ms gate time may be selected and 100 samples taken and applied to the following formula:

$$AVAR = \frac{1}{2(N-1)} \sum_{i=0}^{N} (f(i) - f(i-1))^2 \qquad (4)$$

Where $f(i) - f(i-1)$ is the difference between successive frequency measurements (MIL-PRF-55310D). [7]

The result gives us a feel for how stable the oscillator is reading to reading at the given gate time. By lengthening the gate time the oscillator's Allan variance decreases, showing it is more stable over longer averaging periods.

■ **Jitter**

From this measurement we can see that the instantaneous frequency of the oscillator is not constant – it varies slightly about the nominal frequency creating an uncertainty in the frequency at any given point in time. This frequency change can be viewed as a change in time of the waveform edge from the ideal nominal frequency edge. This change in time of the edge is called jitter. Figure 2 [1] illustrates the jitter effect on a square wave.

Figure 2
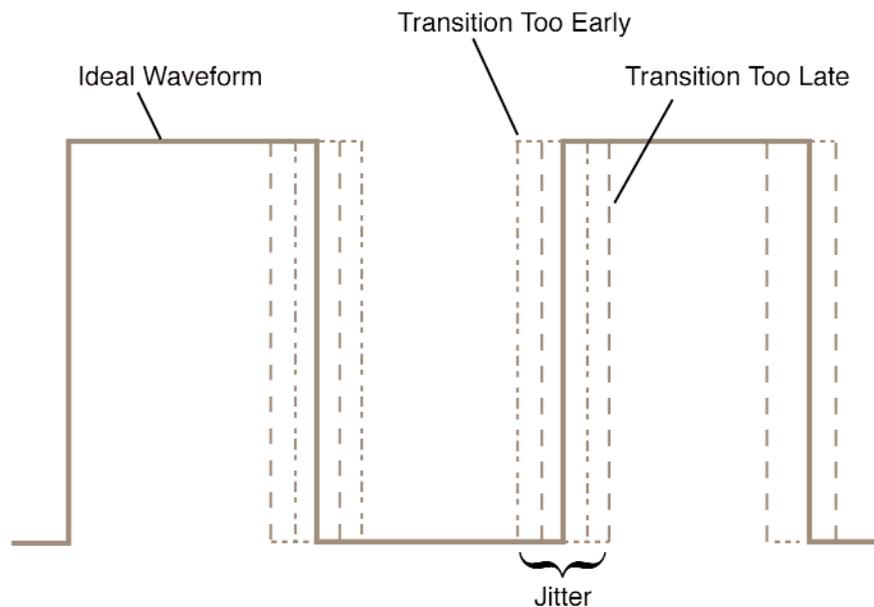
Figure 2: Time domain representation of jitter. [1]

Jitter can be measured in the time domain and is expressed in a peak-to-peak time variation of the edge. This method however may not be very useful in some applications because the variations in the edges are coming from the entire band of frequencies and exaggerating the magnitude of the jitter. Most real world applications will operate within a certain band of frequencies and therefore the jitter effect only needs to be measured in that band. To effectively see and measure jitter over a particular frequency band, conversion to the frequency domain must take place.

■ Phase Noise

This conversion to the frequency domain results in a measurement called phase noise. It is expressed as a graph of power vs. frequency. To understand this measurement the jitter to phase noise conversion must be explored. One way to measure jitter is to measure the variance of each period from the average period, as per Equation 5 [5].

$$\sigma_c^2 = \lim_{N \to \infty} \left(\frac{1}{N} \sum_{n=1}^{N} (\tau_n - \tau_{avg})^2 \right) \tag{5}$$

This RMS cycle jitter can then be used to calculate the phase noise at a given frequency, as can be seen in Equation 6.

$$\pounds(f) = \frac{\sigma_c^2 \, f_{osc}^3}{f^2} \tag{6}$$

Where f_{OSC} is the frequency of the oscillator and f is the frequency away from the carrier [5]

This computation can be done at many discrete frequencies and compiled into graphical form. The above formulas [5, 6] are assuming that there is no 1/f noise or burst noise [5]. In real world oscillators these noise sources are present and shall be discussed below.

If we take into account real world components and the noise that is generated in these circuits, the phase noise calculation gets more complex. The Leeson equation, as shown below in Equation 7, gives us a feel for how circuit noise and circuit elements factor into the phase noise measurement [6].

$$\pounds(f_m) = 10\log\left\{ \frac{FkT}{2P_{avs}} \left[1 + \frac{f_c}{f_m} + \left(\frac{f_o^2}{2f_m Q_t} \right)^2 \left(1 + \frac{f_c}{f_m} \right) \right] \right\} \tag{7} \quad [10]$$

Where Q_1 is the loaded Q of the circuit, f_m is the frequency from the carrier, f_c is the flicker noise corner frequency, f_o is the carrier (oscillator) frequency, T is the temperature in Kelvin, P_{avs} is the power through the resonator, F is the noise factor of the active device, and k is the Boltzmann constant.

Figure 3 (next page) shows how this equation fits into the plot of phase noise. It can be seen that close into the carrier, flicker noise dominates the curve and has a cutoff frequency of the corner frequency of the active device. The middle portion of the phase noise plot follows Leeson's equation and is a combination of loaded Q, noise factor, power and temperature [6]. For frequencies above $f_o/(2Q_l)$ the floor is determined by noise factor, temperature and power.

Figure 3
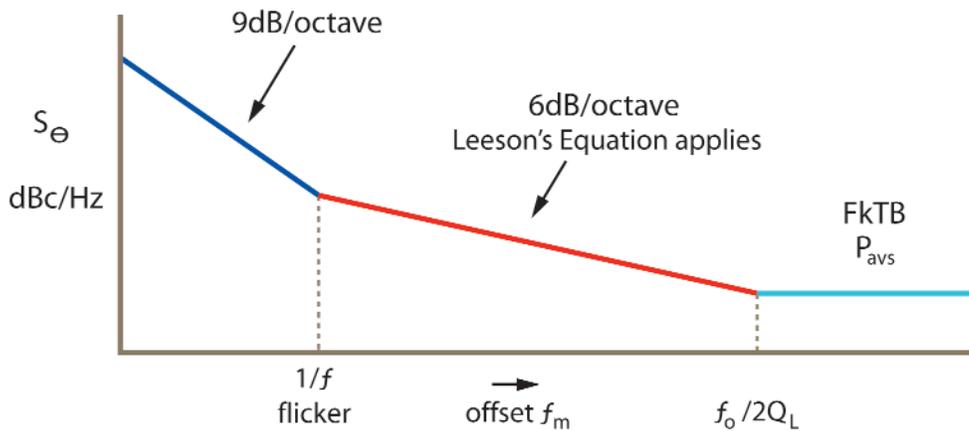
Fig. 3: Phase noise plot showing determining factors [6]

From this plot it is fairly easy to come up with guide lines for minimizing phase noise in oscillator designs. Use devices with low flicker noise. Since the 9dB/octave section is dominated by this quantity, reduction of circuit flicker noise is of great concern. BJTs have a much lower flicker noise than FETs making them more suitable for low phase noise applications. The 6dB/octave section implies that Q of the circuit being higher is of great interest as is noise factor and power. As will be described below, OCXOs utilize crystals with higher Qs than TCXOs. Higher drive power is also desirable because that is the driving factor for the phase noise floor (frequencies above $f_o/(2Q_l)$). This comes with a trade-off because higher drive levels usually result in some phase noise degradation close in to the carrier.

■ Phase Noise Measurement

Measuring phase noise is not an easy feat. Most spectrum analyzers do not have the resolution to measure the phase noise of a crystal oscillator directly. Figure 4 shows the normal configuration for phase noise testing. An 'ideal' source that is the same frequency as the oscillator under test is mixed against the oscillator. This produces the products of these two signals and also the difference. Using a low pass filter, the products are stripped away leaving only the difference, which will be the oscillator noise, if the 'ideal' source is exactly the same frequency as the oscillator.

Figure 4
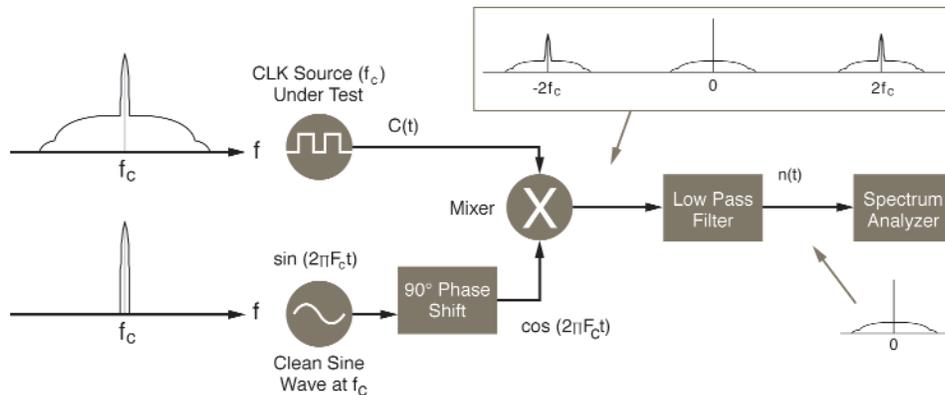
Figure 4. Phase Noise Setup [2]

T his resulting noise signal can then be measured by a spectrum analyzer and displayed in graphical form as a graph of dBc/Hz vs. Hz. This graph is showing us, in terms of power, how far the noise (at each particular frequency) is below the carrier (oscillator's desired signal).

Figures 5 and 6 (below) are typical phase noise measurement plots. Figure 5 shows an OCXO (2 inch square package) plot which has a floor of around -170dBc and a 10 Hz performance of approximately -130dBc. Figure 6 depicts the phase noise performance of a 5mmx7mm TCXO, which has a floor of around -155dBc and a 10 HZ performance of approximately -90dBc.

Figure 5
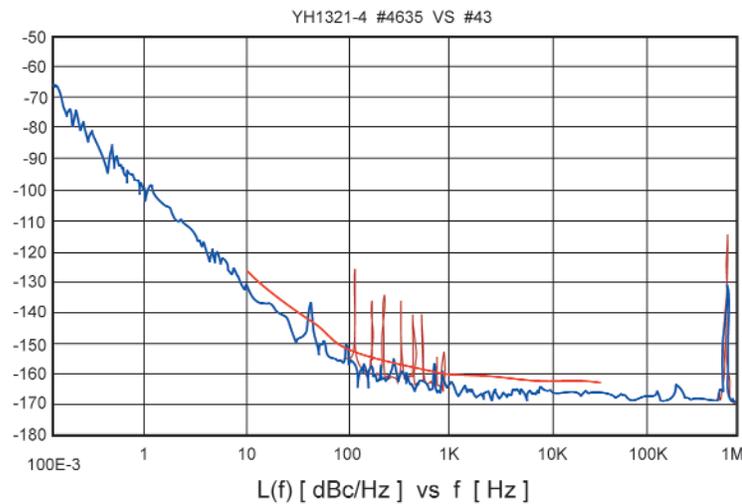
Figure 5. Greenray Industries YH1321-4 OCXO Phase Noise Performance

Figure 6
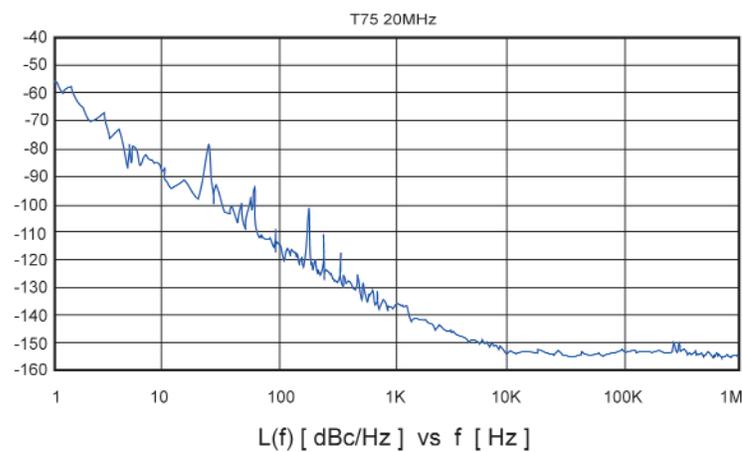
Figure 6. Greenray Industries T75 TCXO Phase Noise Performance

OCXO phase noise performance is commonly much better than that of TCXOs. This is due to several factors. The OCXO is inherently more stable than the TCXO giving it improved phase noise performance. This is because the crystal in the OCXO is thermally stable where the TCXO's crystal is feeling the environmental changes, and even though there is compensation for those ambient temperature changes it is nowhere near as stable as the OCXO.

Secondly, OCXOs use different crystal cuts than TCXOs and as a result have a higher Q. A TCXO's crystal Q is approximately 30,000-40,000 where an SC cut ovenized crystal may have a Q approaching 1,000,000. This higher Q directly helps improve phase noise in the oscillator.

Lastly, TCXOs are designed to be much smaller and have very low power consumption. This means most of them are constructed using FETs instead of BJTs. FETs have a much higher inherent flicker noise which worsens the phase noise characteristics of the oscillator. Ovenized units designed for stability use BJTs which have lower noise specifications and in turn better phase noise performance.

■ Measurement Flaws

One of the issues with this measurement system is that you are mixing the oscillator under test against an "ideal, noise free" oscillator. Since there is no such real world device you must either mix it against something you know is much better than the unit under test and assume that the 'reference' injects no noise. Another option is to mix against something that is equivalent to the unit under test, and assume equal contribution from both oscillators. Both of these methods are inherently flawed. By mixing against something better you are assuming no degradation from the reference.

If you do have a reference that is 10 times better the effects are minimal but there is still degradation in your measurement. By mixing against an equal source and assuming equal contribution you add 3dB to your target, which allows you to accept results which really exceed the intended specification because there is another noise source. However, how do you really know that they are equal, and in fact most times they are not equal. One source will have better noise characteristics than the other and this creates ambiguous data. In fact if the unit picked as a 'reference' ends up being a relatively clean source, and you assume equal contribution you could be accepting results that are in reality failing, but you have no way of knowing.

■ Outside Noise/Interference

Aside from the circuit characteristics, measurement assumptions, and random crystal fluctuations, the oscillator and its phase noise behavior are affected greatly by the power supplied to the oscillator. Figure 7 (see next page) shows a TCXO with a noisy power supply. Spurs can be seen between 80 and 1000 Hz with an amplitude of approximately 40 dBc.

Fig. 8 (see next page) shows the same TCXO tested with a cleaner supply and only two spurs remain with amplitudes of approximately 20 dBc. The oscillator manufacturer has no control over what power supply the customer uses in their application and therefore attempts to characterize just the phase noise performance and not the impact of the supply on the oscillator. Oscillator manufacturers try to use the cleanest supplies possible for these measurements.

Figure 7
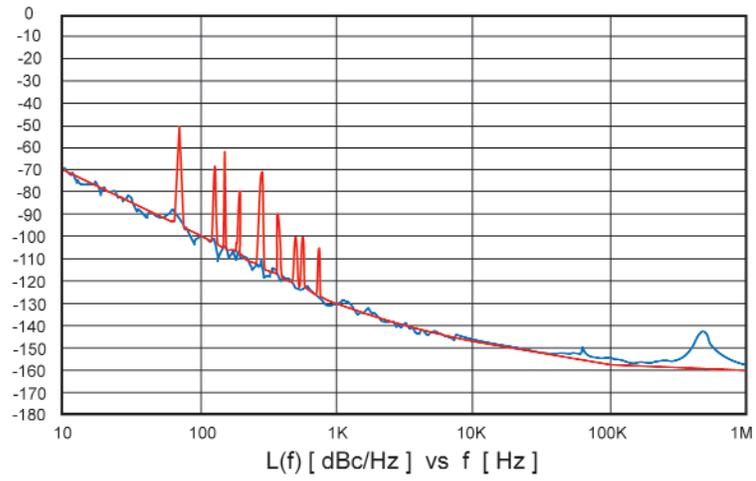
Figure 7. TCXO with noisy power supply

Figure 8
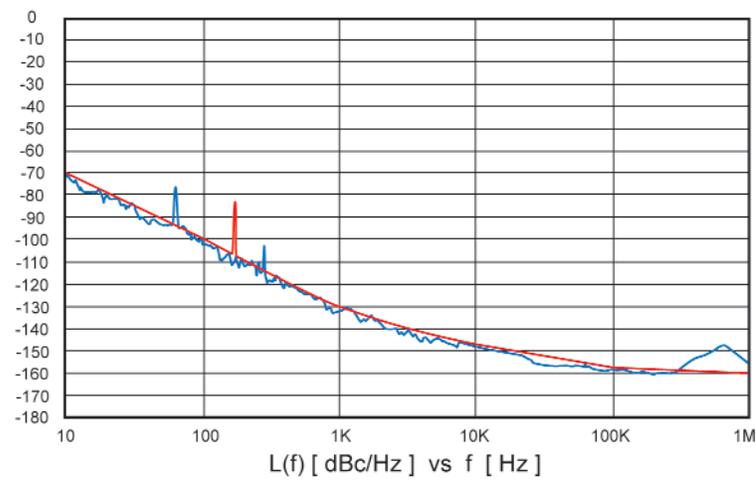
Figure 8. TCXO with cleaner power supply

Nearby interference can also show up in phase noise measurements. Fig. 9 (below) shows the TCXO from above with a clean supply running it, but a computer monitor running about 3 feet away from the oscillator and test set. Fig. 10 shows the same test with the monitor turned off. Again the spurs from Figure 9 are not from the oscillator and therefore we attempt to eliminate them by shutting off the computer monitor during testing.

Figure 9
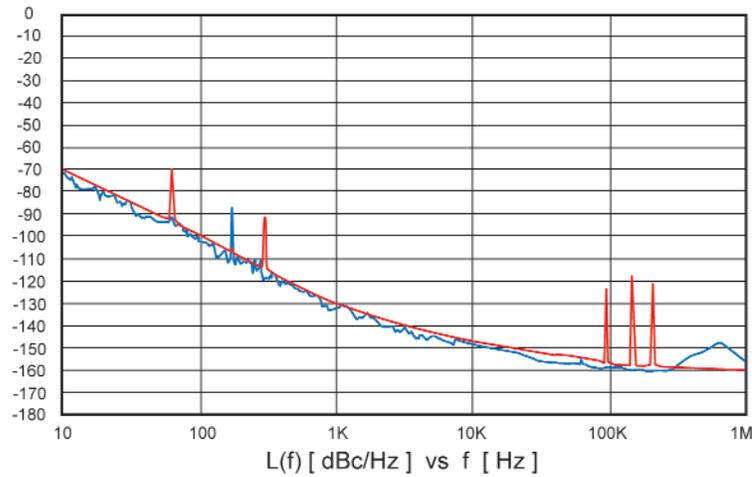
Figure 9. TCXO with monitor running

Figure 10
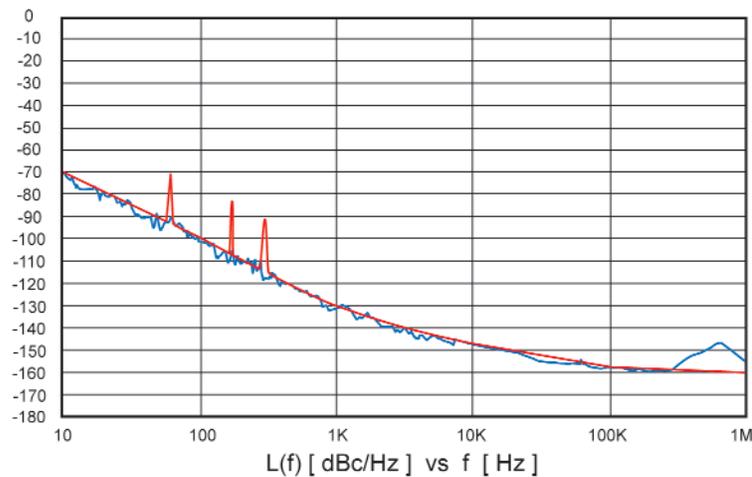
Figure 10. TCXO with monitor off

These real world plots (Figures 9, 10) show how a potential disconnect between consumer and manufacturer can occur.

scillator manufacturers characterize just the oscillator's phase noise performance and may go to great lengths to demonstrate "low" phase noise. Clean power supplies or even batteries are used in an attempt to reduce extraneous noise; nearby interference sources such as computer monitors or other equipment are mitigated. Manufacturers may even build a Faraday Cage - a shielded room - to reduce outside interference. Many use baffles or buffers to reduce airflow around the oscillator to improve thermal conditions and reduce phase noise.

Unfortunately, design engineers specing/purchasing oscillators may not know that these measures were taken and may not realize the ultimate impact for their particular system or application. The oscillator will not be running in a shielded room all by itself. It will be in a larger assembly with temperature gradients and non-ideal power supplies and other machinery or computers running nearby.

When a design engineer looks at a phase noise plot and assumes that this is what they have in their system, there could be a problem. If no headroom is allowed for in their design, serious problems could occur.

Lastly, there are two additional factors that directly effect phase noise performance which oscillator manufacturer have no control over. Vibration affects the phase noise of crystal oscillators as in the following formula:

$$\pounds(f_v) = 20\log(\frac{\Gamma A f_o}{2f_v})$$
(8)

Where _ is the oscillator's g-sensitivity, A is the amount of g's in the vibration, f_o is the frequency of the oscillator, and f_v is the frequency of the vibration. [10]

Most military designers are well aware of this degradation and work with the manufacturer to understand and characterize the oscillator under the vibration it will see in the field.

Another factor that hinders phase noise performance is multiplication. Many customers will buy an oscillator in the megahertz and multiply it up into the Gigahertz range. This is acceptable and will be a stable oscillator but phase noise is degraded by:

$$\pounds(f) = 20\log(N)$$
(10)

Where N is the factor of multiplication. [10]

Vibration and multiplication are two processes that the user might subject the oscillator to, and potential performance degradation should be addressed with the oscillator manufacturer. It is important for designers to understand the effects all these variables can have on the time base - and to work with the manufacturer to ensure delivery of a quality oscillator that satisfies their real world, performance requirements.

■ **Conclusion**

Phase noise is a very important aspect of any time base. Oscillator manufacturers seek to minimize phase noise in an oscillator by maximizing Q, using low noise active devices, and by increasing resonator power.

The method of testing phase noise directly effects test results: low noise supplies or batteries may be used to reduce phase noise from supply ripple; nearby interference from computers, lighting, or other machinery may be eliminated. Finally, thermal buffers may be used to ensure that the test is only measuring the oscillator noise and not noise from outside sources.

For systems design engineers using oscillators as a time base, understanding the phase noise phenomenon and characterizing its effects are essential. System designers should communicate with the oscillator manufacturer to determine the likelihood and the impact of oscillator phase noise on their system design and ultimate, real world performance.

■ Looking Forward

Oscillator designers are always looking for ways to improve phase noise performance and provide a more stable oscillator to the customer. Many noise reduction schemes are being attempted and implemented. Currently Greenray Industries is making a TCXO that can achieve better than -100 dBc at 10Hz by using close in filtering on the compensation voltage. Other ideas to be tested would be using a crystal filter in an OCXO to act as a super-selective band pass filter. This should reduce the phase noise floor but several factors must be overcome for practical implementation.

The future is indeed bright for high precision, low-noise oscillators and their critical role in a wide variety of demanding applications.

Works Cited:

[1] Neil Roberts, "Phase Noise and Jitter – A Primer for Digital Designers", EE Design, July 14, 2003
[2] "Clock Jitter and Phase Noise Conversion", Maxim Dallas Semiconductor, 2004
[3] Steve Fry, "Quartz Crystal Based Low Phase Noise Oscillator Characteristics", Greenray Industries, 2005
[4] Barry Kleinle, "Phase Noise in Crystal Oscillators", Greenray Industries, 2004
[5] Rick Poore, "Phase Noise and Jitter", Agilent, 2001
[6] "Phase Noise", [Online] - http://www.zen118213.zen.co.uk/Systems_And_Devices_Files/PhaseNoise.pdf
[7] Oscillator, Crystal Controlled General Specification For, MIL-PRF-55310D, 1998
[8] Dr. Lee A. Breakiron, Cesium Atoms at Work, http://tycho.usno.navy.mil/cesium.html
[9] NIST-F1 Cesium Fountain Atomic Clock, http://tf.nist.gov/timefreq/cesium/fountain.htm
[10] John R. Vig, "Quartz Crystal Resonators and Oscillators", U.S. Army Electronics Technology and Devices, Fort Monmouth, NJ, 1990

Author: John Esterline, Greenray Industries, Inc.

This article was originally published in Microwave Journal, April 2008

12
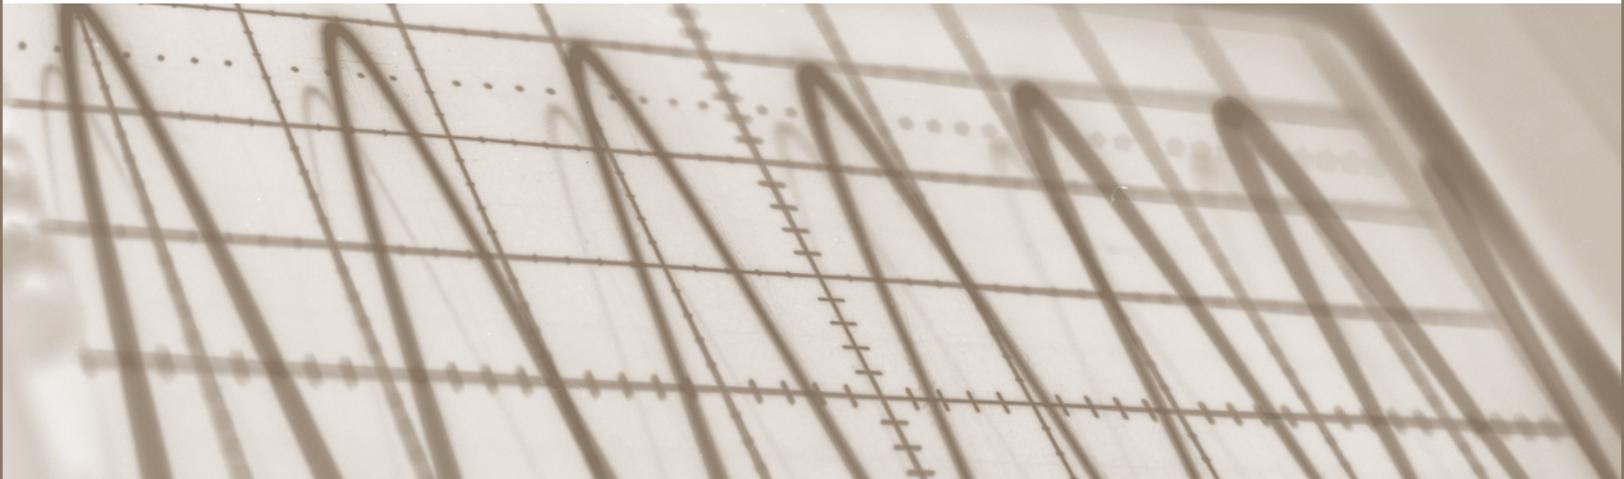
Precision quartz-based oscillators for the wireless, wired telephony, aerospace, military, satellite, and other communications markets.

GREENRAY