VARIANCE AS APPLIED TO CRYSTAL OSCILLATORS

Before we can discuss VARIANCE AS APPLIED TO CRYSTAL OSCILLATORS we need to understand what a Variance is, or is trying to achieve. In simple terms a Variance tries to put a meaningful figure to 'what we actually receive' against 'what we expect to receive'. It is, simply, a mathematical formula applied to a set of data points / samples / readings which are usually collected over a specified period of time. There are various types of Variances, each tailored to suit a particular application. Variance is σ_y^2 , but the term Variance is also used for $\sqrt{\sigma_y^2}$, it is up to the individual to interpret which is being quoted. Variance is only useful if it converges. Convergence means the more samples we take the closer the resulting Variance gets to a steady value. Non convergence means the Variance just gets bigger and bigger as we take more and more samples.

An example of a converging Variance is the number of times the flip of a regular coin turns up heads. The more times we flip the coin the more likely the Variance is to converge to 0.25 (0.25 is 0.5^2). An example of a non converging Variance is the age of a person, the more data we collect the larger the Variance gets, it is not heading for a steady value. This means we have to have an understanding of the underlying causes of the variability of the collected data before we can decide if a Variance will have any meaning.

For a crystal oscillator 'what we expect to receive' is a fixed / stable frequency that never changes. 'What we actually receive' is very close to a fixed / stable frequency but it is perturbed by the crystal oscillators inherent noise sources. This means Variance is another way of measuring the stability of a crystal oscillator in the time domain, (as is jitter). To understand the underlying inherent noise sources of a crystal oscillator it is useful to consider the stability of the crystal oscillator in the frequency domain, i.e. the crystal oscillators Phase Noise. These inherent noise sources are covered by the article PHASE NOISE IN CRYSTAL OSCILLATORS.

Variance, Jitter and Phase Noise are all inter related, the choice of which to use when considering the stability of a crystal oscillator is usually application specific. RF (Radio Frequency) Engineers working in Radar or Base Station design will be interested in Phase Noise as poor Phase Noise performance will affect Up/Down conversions and channel spacing. Digital Engineers working in Time Division Multiplexing (the majority of modern Telecoms infrastructure) will be interested in jitter as poor jitter performance will result in Network slips and excessive re-send traffic. Engineers working in GPS will be interested in Variance as poor Variance can increase acquisition lock times and may cause loss of lock. Each application is interested in a different part of the Phase Noise spectrum.

As stated earlier there are various types of Variances, each tailored to suit a particular application. As this discussion is about VARIANCE AS APPLIED TO CRYSTAL OSCILLATORS we will consider Standard Variance, and show why it is not suitable for crystal oscillator stability measurements, Allan Variance and Hadamard Variance, which are suitable for crystal oscillator stability measurements. In particular we will consider the frequency stability of the oscillator with respect to time.

Geoff Trudgen, Rakon UK Ltd, July 2009

Standard Variance, for a population of samples (the data), is the mean (arithmetic average) of the squares of the differences between the respective samples and their mean. It attempts to put a single value to the extent the population of samples (the data) varies from the 'average' value. If the Variance σ_y^2 is close to zero then the population of samples (the data) is closely packed. A large Variance σ_y^2 says the population of samples (the data) is widely spread out. This makes Variance dimensionless.

Mathematically it is expressed as (Fig. 1).

Figure 1

$$\sigma_y^2 = \frac{1}{m-1}\sum_{i=1}^{m}(y_i - \bar{y})^2$$

Where:-

m	is the number of samples
y_i	is the value of sample i
$\bar{y}$	is the mean (arithmetic average) of the samples
$\sigma_y^{\ 2}$	is the Variance

Note:- The square root of the Variance is the Standard Deviation.

Consider these three sets of data (Fig. 2).

Figure 2

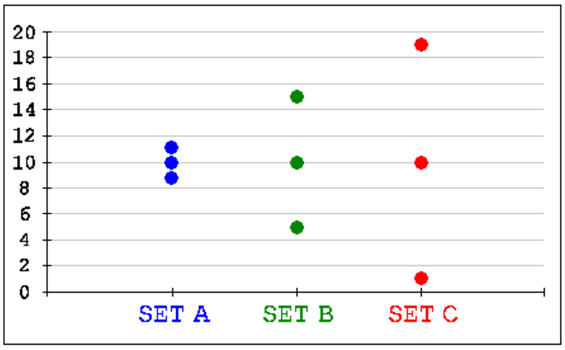

Set A 9, 10, 11
Set B 5, 10, 15
Set C 1, 10, 19

The mean (arithmetic average) $\bar{y}$ of the samples is 10 for all three sets.

Set A $\bar{y}$ = (9+10+11)/3 = 10
Set B $\bar{y}$ = (5+10+15)/3 = 10
Set C $\bar{y}$ = (1+10+19)/3 = 10

But the Variance σ_y^2 is significantly different.

Set A σ_y^2 = [(9-10)^2 + (10-10)^2 + (11-10)^2] / (3-1) = 1
Set B σ_y^2 = [(5-10)^2 + (10-10)^2 + (15-10)^2] / (3-1) = 25
Set C σ_y^2 = [(1-10)^2 + (10-10)^2 + (19-10)^2] / (3-1) = 81

Showing the spread of data Set A ,with a low value of Standard Variance, is significantly tighter than the spread of data Set C ,with a high value of Standard Variance. Remember Variance is only useful if it converges. Standard Variance will only converge for a sample set that has a Gaussian type distribution, (the actual length of a sample of 50mm M6 screws for example). A sample set with a systematic drift or discontinuities (jumps in the data) will not converge.

When applied to a crystal oscillator we need to consider the different types of inherent noise sources within the oscillator. Fig.3 is a Spectral Density Plot (Idealised Phase Noise Plot) showing the various noise types for an oscillator. Standard Variance will only converge for White Phase Modulation (f), Flicker Phase Modulation ($1/f^1$) and White Frequency Modulation ($1/f^2$), the noise sources with a Gaussian type distribution. It will not converge for higher orders of noise, Flicker Frequency Modulation ($1/f^3$) and higher. This is why Standard Variance is not suitable for measuring a crystal oscillators frequency stability over time. The article PHASE NOISE IN CRYSTAL OSCILLATORS explains the concept of Spectral Density.

Figure 3

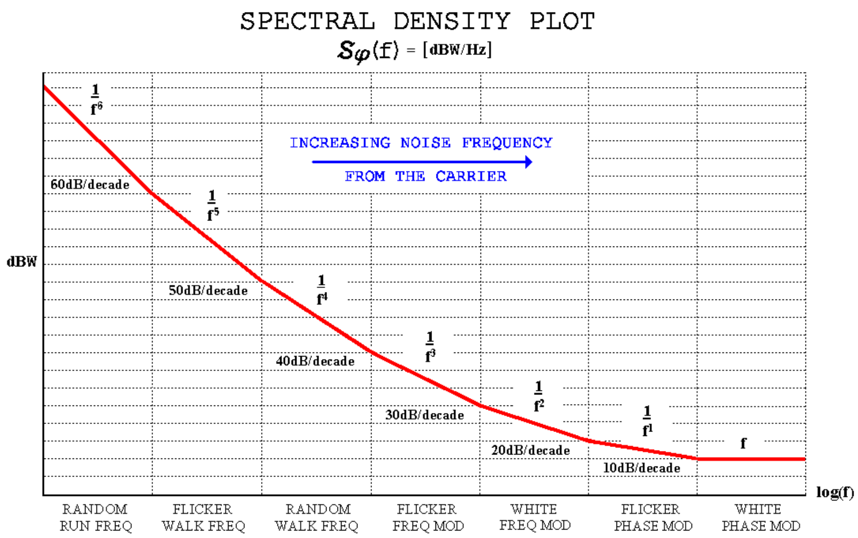

Note:- This plot is a log / log plot (dBW are logarithmic).

For a precision TCXO (as a very approximate guide) these noise sources cover.

Noise Source	Slope	Offset Frequency
White Phase Modulation	(f)	>10kHz
Flicker Phase Modulation	($1/f^1$)	1kHz to 10kHz
White Frequency Modulation	($1/f^2$)	10Hz to 1kHz
Flicker Frequency Modulation	($1/f^3$)	100mHz to 10Hz
Random Walk of Frequency	($1/f^4$)	1mHz to 100mHz
Flicker Walk of Frequency	($1/f^5$)	10uHz to 1mHz
Random Run of Frequency	($1/f^6$)	<10uHz

Allan Variance (also known as **Two-Sample Variance**) was designed to overcome this non convergence. It is one half the mean (arithmetic average) of the squares of the differences between successive frequency readings sampled over a chosen measurement period. The samples must be taken with no dead-time between them. (Dead time between the readings will skew the result). Using the difference between successive frequency readings effectively applies a one pole high pass filter function to the measurements.

Mathematically it is expressed as (Fig. 4)

Figure 4

$$\sigma_y^2(\tau) = \frac{1}{2(m-1)} \sum_{i=1}^{m-1} (y_{i+1} - y_i)^2$$

Where:-

m	is the number of samples
y_i	is the value of sample i
y_{i+1}	is the value of sample i+1
τ	is Tau, the sample time
$\sigma_y^2(\tau)$	is the Allan Variance

Note:- The divide by two causes Allan Variance to be equal to the Standard Variance if y_i comes from a random and uncorrelated set, i.e. white (Gaussian) noise.

The advantage of Allan Variance, when applied to Crystal Oscillators, over the Standard Variance is it will also converge for Flicker Frequency Modulation ($1/f^3$) and Random Walk of Frequency($1/f^4$). It does not converge for Flicker Walk of Frequency ($1/f^5$) and Random Run of Frequency ($1/f^6$).

The Allan Variance for a crystal oscillator is usually quoted as the Root Allan Variance (RAV), i.e. the square root of $\sigma_y^2(\tau)$. For a precision crystal oscillator it is typically a number around 70×10^{-12} (70 pico). Occasionally the RAV is quoted in Hz which is simply generated by multiplying the RAV by the Oscillator frequency. I.e. a 10MHz oscillator with a RAV quoted as 0.68mHz actually has a RAV of 68 pico (0.68mHz/10MHz).

To measure RAV requires a frequency counter with the ability to take continuous frequency readings with no dead time between the measurements. One such counter is the Pendulum CNT-90. This frequency counter must be locked to an ultra stable frequency standard with excellent RAV. A good reference oscillator for the frequency standard is the RAKON CFPODO 10MHz OCXO. The frequency counter has to accurately measure frequencies to sub 0.1 parts per billion (0.1×10^{-9}) so it is usual to employ a Heterodyne method. The synthesiser used for Heterodyning must also be low phase noise and locked to the same frequency standard as the counter.

A gate time needs to be selected for the frequency counter, and the number of readings to use needs to be decided before a RAV measurement can be made. The combination of gate time and number of readings (the total sample time) must not be so large as to include the Flicker Walk of Frequency noise ($1/f^5$).

A 1 second gate and 100 readings (100 seconds equates to an Offset Frequency of 10mHz) covers the White Phase Modulation (f) to Random Walk of Frequency($1/f^4$) noise sources and gives a good indication of the RAV for a precision crystal oscillator.

Fig.5 is the same Spectral Density Plot (Idealised Phase Noise Plot) but shown in the Time Domain as a Variance Tau Plot.

Figure 5

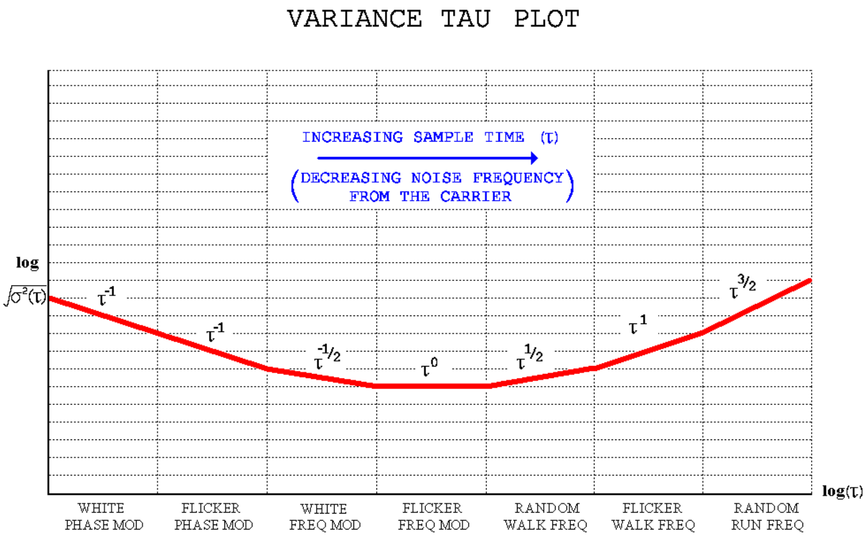

Note:- This plot is a log / log plot (Noise source legends are shown reversed compared to the Spectral Density Plot).

Using the same approximate frequencies as the Spectral Density Plot, these noise sources cover.

Noise Source	Slope	Sample Time (Tau)
White Phase Modulation	(τ^{-1})	<10µs
Flicker Phase Modulation	(τ^{-1})	10µs to 1ms
White Frequency Modulation	$(\tau^{-1/2})$	1ms to 100ms
Flicker Frequency Modulation	(τ^{0})	100ms to 10s
Random Walk of Frequency	$(\tau^{1/2})$	10s 1ks
Flicker Walk of Frequency	(τ^{1})	1ks to 100ks
Random Run of Frequency	$(\tau^{3/2})$	<100ks

Note:- 1ks (kilo second) is ~17 mins. 100ks is ~28 hours.

Allan Variance does not distinguish between White Phase Modulation (f) and Flicker Phase Modulation ($1/f^1$), hence the slope for both types of noise sources is τ^{-1}.

Allan Variance will not converge for Flicker Walk of Frequency ($1/f^5$) and Random Run of Frequency ($1/f^6$)

Hadamard Variance (also known as **Three-Sample Variance**) was designed to overcome the non convergence of Allan Variance. It is one sixth the mean (arithmetic average) of the squares of the differences between three successive frequency readings sampled over a chosen measurement period. The samples must be taken with no dead-time between them. Using the difference between three successive frequency readings effectively applies a 2 pole high pass filter function to the measurements.

Mathematically it is expressed as (Fig. 6).

Figure 6

$$\sigma_y^2(\tau) = \frac{1}{6(m-2)} \sum_{i=1}^{m-2} \left(y_{i+2} - 2y_{i+1} + y_i \right)^2$$

Where:-

m	is the number of samples
y_i	is the value of sample i
y_{i+1}	is the value of sample i+1
y_{i+2}	is the value of sample i+2
τ	is Tau, the sample time
$\sigma_y^2(\tau)$	is the Hadamard Variance

When applied to crystal oscillators the advantage of the Hadamard Variance over the Allan Variance is for longer sample times as the Hadamard Variance will also converge for Flicker Walk of Frequency ($1/f^5$) and Random Run of Frequency ($1/f^6$). The Hadamard Variance is unaffected by linear frequency drift, effectively removing the Oscillator ageing effect from the Variance measurement for long sample times.

For the noise sources White Phase Modulation (f) to Flicker Frequency Modulation ($1/f^3$) the Allan Variance and Hadamard Variance are comparable (Fig. 7) but for Random Walk of Frequency ($1/f^4$) the Hadamard Variance is ~half the Allan Variance.

Figure 7

Noise Source	Slope	Allan Variance	:	Hadamard Variance
White Phase Modulation	(f)	1	:	0.9
Flicker Phase Modulation	($1/f^1$)	1	:	1.0
White Frequency Modulation	($1/f^2$)	1	:	1.0
Flicker Frequency Modulation	($1/f^3$)	1	:	1.2
Random Walk of Frequency	($1/f^4$)	1	:	0.5

Next step:

For more information email: info@rakon.co.uk

Geoff Trudgen, Rakon UK Ltd, July 2009