


Measurement Good Practice Guide

Good Practice Guide to Phase Noise Measurement

David Owen
Pickering Interfaces

No. 68



Measurement Good Practice Guide No. 68

Good Practice Guide to Phase Noise Measurement

David Owen
Pickering Interfaces

Abstract: Measurement of phase noise has become important in designing and developing systems for communications which pervade so much of our daily life. Without good, low phase noise oscillators in systems such as mobile telephone systems, the levels of interference experienced by users would become unacceptable. Improvements in these systems are often contingent upon improved oscillators and these improvements can only be quantified through good measurements of phase noise.

© Crown Copyright 2004
Reproduced by permission of the Controller of HMSO

ISSN 1368-6550

May 2004

National Physical Laboratory
Teddington, Middlesex, United Kingdom, TW11 0LW

Website: www.npl.co.uk

Acknowledgements

This guide has been written by David Owen on behalf of the National Physical Laboratory.

David Owen was involved in the measurement of phase noise on VCO's, synthesizers and crystal oscillators for many years while employed by Marconi Instruments, which became IFR, and is now part of Aeroflex. David was first involved with phase noise while working as a design engineer of signal generator and signal analyzer products, and later as the business manager for the signal generator group. David Owen is currently the PXI Business Development Manager for Pickering Interfaces.

Editorial support for this guide was provided by David Adamson of NPL.

A Good Practice Guide to Phase Noise Measurement

Contents

| | | |
|-------|--|----|
| 1 | Introduction to phase noise | 1 |
| 1.1 | Phase noise and frequency noise | 3 |
| 1.2 | Scaling noise and bandwidth | 4 |
| 1.3 | Single sideband and double sideband | 5 |
| 1.4 | Amplitude noise | 5 |
| 1.5 | Comparing phase noise profiles | 6 |
| 1.6 | Relevance of phase noise to communication systems | 8 |
| 1.6.1 | Analog communication systems | 9 |
| 1.6.2 | Digital communications | 9 |
| 2 | General principles of methods of measuring phase noise | 10 |
| 3 | Spectrum analyzer methods | 11 |
| 3.1 | Down converter and IF filter with a spectrum analyzer | 14 |
| 3.2 | Spectrum analyzer measurement calibration | 16 |
| 4 | Delay line discriminator | 17 |
| 4.1 | Calibration of delay line discriminators | 19 |
| 5 | Quadrature method | 21 |
| 5.1 | Calibration of the quadrature method | 24 |
| 5.2 | Calibration of higher offsets | 24 |
| 5.3 | Calibration of PLL effects (close to carrier phase noise measurements) | 26 |
| 6 | FM discriminator method | 28 |
| 6.1 | Calibration of FM discriminator | 30 |
| 7 | Digitizer Measurements | 31 |
| 7.1 | Calibration of Direct Digital Measurement | 33 |
| 8 | Common problems with phase noise measurement | 34 |
| 8.1 | Reference oscillator | 34 |
| 8.2 | Injection locking | 34 |
| 8.3 | Reference oscillator locking | 35 |
| 8.4 | Mixer linearity and errors | 35 |
| 8.5 | Power Supplies | 37 |
| 8.6 | Detectors and filters | 38 |
| 8.7 | Stitching errors | 39 |
| 8.8 | Spurious noise identification and display | 39 |
| 8.9 | Dithered clocks | 40 |
| 8.10 | Software correction artifacts | 40 |
| 8.11 | Sideband selectivity | 40 |
| 8.12 | AM Noise rejection | 41 |
| 9 | Comparison of the Methods | 43 |
| 10 | Estimating Uncertainty | 48 |
| 10.1 | Spectrum Analyzer Uncertainty | 48 |
| 10.2 | System Level Sensitivity | 52 |
| 10.3 | PLL Estimation | 53 |
| 10.4 | Using an FM signal generator for calibration | 53 |
| 10.5 | Relative Accuracy of Systems | 54 |
| 11 | Other Measurements | 55 |

| | | |
|------|---|----|
| 11.1 | Two Port Phase Noise Measurements..... | 55 |
| 11.2 | Calibration of Two Port Measurement..... | 57 |
| 11.3 | Measurement of Phase Noise on Pulsed Signals..... | 57 |
| 12 | References | 59 |
| 13 | Glossary..... | 60 |

1 Introduction to phase noise

A perfect frequency source generates only one output signal with no instability in its output frequency or level. Both the output level and frequency are measurable to great precision. In reality, however, all signal sources exhibit some instability in both their instantaneous output frequency and their level. This guide is primarily intended to show how to measure the short-term frequency (phase) instability, but some references will also be made to the measurement of level (amplitude) uncertainty since in some circumstances they can be measured by the same test systems.

The short-term frequency instability can be expressed in a number of different ways. The method of expressing the instability is likely to depend upon the intended application, as well as the performance of the signal source, and in many cases a source may be characterised in more than one way. While the measurement methods used are in principal relatively easily understood, there are a surprising number of complications to the measurement if good accuracy and consistent measurements are required (for example when trying to improve the performance of an oscillator where small changes may be significant).

Phase noise is the most generic method of expressing frequency instability. The carrier frequency instability is expressed by deriving the average carrier frequency and then measuring the power at various offsets from the carrier frequency in a defined bandwidth. The result is then expressed as a logarithmic ratio compared to the total carrier power. The power ratio is usually normalised to be the equivalent signal power present in a measurement bandwidth of 1 Hz. For some applications (e.g., specifying adjacent channel power on a transmitter), it can be expressed in other bandwidths (in the case of adjacent channel power the receiver bandwidth is usually used).

It should be noted that phase noise measurements are inherently ratiometric in nature. In principle, the absolute level of the carrier signal to be measured is not relevant. The measured performance of the source should not change with the level it is measured at, unless the source operating conditions are changed (the addition of an attenuator pad for example should not fundamentally change the performance).

The fact that phase noise is normally expressed in terms of the relative power (to the total carrier power) in a 1 Hz BW does not mean that the signal is actually measured in a 1 Hz bandwidth. This can lead to some confusion, for example:

- Coherent signals on a phase noise plot may need to have different correction factors applied to them compared to noise signals. The coherent signal occupies a single 1 Hz frequency bin while the noise signal is spread out over all the frequency bandwidth of the measurement system
- At close to carrier frequencies the power referenced to 1 Hz could be higher than the total carrier power. In this case, the phase noise signal is measured in a bandwidth much less than 1 Hz and then has to be “normalized” to a 1 Hz bandwidth.

The phase noise on a source can be caused by a number of different factors. These include (but are by no means limited to)

- Noise in the semiconductor devices that generate the signal
- Noise from power supplies that is converted to a phase deviation by non-linear processes
- Varactor diodes internal noise (or other frequency tuning devices)
- Noise on the tuning lines to frequency tuning elements
- Additive noise from amplifier systems, including white noise
- Phase locked loops in synthesized sources

Some of these mechanisms generate phase noise by directly phase modulating the signal; others generate noise in less direct ways.

There are other ways of expressing the phase noise of a source that are more application orientated. The most useful general measurement, however, is the phase noise characteristics since, from a phase noise plot the other measurements can be estimated.

Data sheets for various types of product show that phase noise can vary significantly between different types of devices, and some manufacturers will show actual phase noise profiles while others will show stylized “cleaned up” versions without the spurious signals that normally show (and sometimes confuse) the plots.

The examples below show a stylized plot and a test system measurement for real low noise signal generators¹.

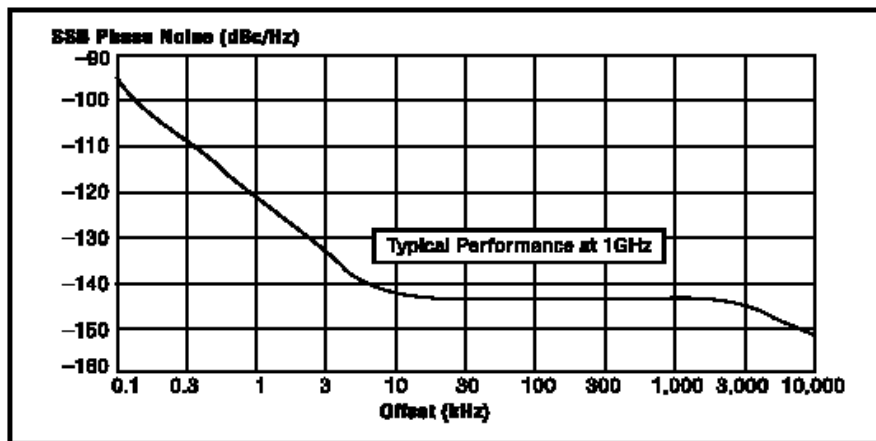


Figure 1: Phase noise plot excluding spurious signals

¹ Source. Aeroflex 2040 series data sheet and Agilent 8644 data sheet




Figure 2: Phase noise plot including spurious signals

The upper trace shows a plot of phase noise which has excluded spurious signals, while the lower trace shows a response that includes low frequency spurious signals. Note in the lower trace the sudden change in the smoothness of the curve is probably due to a change in the measurement bandwidth and the number of measurement averages used by the test system.

1.1 Phase noise and frequency noise

There is often confusion between the phase noise and frequency noise and, for those not familiar with the differences, relating the two may not be obvious.

In phase noise measurements, the signal at a given offset is measured as the relative power of the signal in a normalized 1 Hz bandwidth compared to the total signal power. This is slightly confusing because, at first sight, this is not a measurement of phase noise at all – there are no radians or degree units in the terms. However, if a signal was disturbed by a phase modulation signal of (say) 0.01 radians at a variable offset between (say) 1 kHz and 100 kHz, the sideband signal level seen on a spectrum display would have a fixed relationship with the carrier level (as it happens in this case it would be -46 dBc), irrespective of the frequency offset of the signal. Applying the reverse argument the signal level in dBc/Hz is therefore directly related to the amount of phase disturbance the noise signal generates in that narrow band of modulation frequencies. The phase noise at a given offset expressed in dBc/Hz can be directly translated into degrees or radians in a one Hertz bandwidth² independently of the offset frequency.

Note that all sources of noise can be converted to the equivalent phase modulation factor regardless of whether they are caused by direct modulation of the source (i.e. are truly phase modulation) or are added by some other method (e.g. frequency modulation, added white noise).

² Strictly this conversion relies on an assumption that the phase noise has a relatively low phase modulation index (<0.1) in the band of interest. If the phase noise index is higher the issues are significantly more complex since the Bessel functions that can be used to convert phase to sideband amplitude become non linear.

In the same way that phase noise expressed in dBc/Hz can be expressed in terms of phase noise in radians or degrees, the noise can also be expressed as FM noise. For narrow-band modulation, the relationship is relatively easy to understand:

A modulating signal at 1 kHz rate and having a frequency deviation of 1 Hz has a modulation index (FM deviation divided by modulation frequency) of 0.001. The equivalent phase deviation is 0.001 radians. Narrow band modulation theory states that the sideband level is:

$$-6 \text{ dBc} + 20 \log(\text{mod index})$$

where the mod index is the ratio of the deviation (Hz) to the modulation frequency (Hz).

In this case, the sideband signal level will be seen at -66 dBc relative to the carrier.

If the modulating frequency is increased to 10 kHz and the deviation (in Hz) remains the same the modulation index will fall to 0.0001 and the sideband signal level it generates will be -86 dBc.

This highlights an important difference in the frequency behavior of frequency modulation sources compared to phase modulation sources – a white noise source of FM will exhibit a 6 dB per octave reduction in noise as the offset is increased, while a white noise source of phase modulation will exhibit a flat noise profile.

1.2 Scaling noise and bandwidth

Most measurement systems do not measure phase noise in a 1 Hz bandwidth. In addition, if the significance of the presence of phase noise on a system is to be understood, the impact has to be assessed over a band of offset frequencies. It is therefore important to understand how to scale the measurement with bandwidth.

Provided a noise signal is substantially flat with frequency over the band of interest and it contains no coherent signals the noise level can be scaled simply by the formula:

$$\text{Noise in } \Delta F = (\text{Phase Noise in 1 Hz}) + 10 \log \Delta F$$

where ΔF is the bandwidth of the signal in which noise is to be measured and the phase noise is measured in dB/Hz.

If the noise level is required to be quoted in a 3 kHz bandwidth (common for audio noise) the noise level will be approximately 33 dB higher than it is in a 1 Hz bandwidth.

If the noise is not flat with frequency then the noise in the band of interest can be calculated by integrating the curve shape of the noise if the shape is known, or more practically, by splitting the noise up into sub bands that can be considered to be flat and then again summing the powers together.

If the band of interest has coherent spurious signals present, then this power summation does not work for those signals. Instead the signals must be assessed individually, their

contributions calculated, and a judgment made on how these signals are related so that so they are either added in power or voltage.

The above also assumes that noise is being measured in a bandwidth defined by so-called “brick wall” filters. The notional filter is assumed to have perfect selectivity, it passes the signal inside its bandwidth with no attenuation or ripple, and responds to none of the signals just outside the bandwidth. Such filters are not feasible. Instead, the **noise bandwidth** of the filter is quoted, essentially reversing the calculation. The filter is assumed to be measuring a noise signal which is flat with frequency, and the equivalent “brick wall” filter bandwidth which would pass the same noise power is calculated. Some of the noise measured is contributed from outside the filter bandwidth but most is from inside the bandwidth. The noise bandwidth of the filter takes account of the shape of the filter skirts and its pass band ripple. The noise bandwidth of a filter does not have to be the same as the 3 dB (or 6 dB) bandwidth – there is no good reason why it should be.

1.3 Single sideband and double sideband

Phase noise is quoted as the noise on one side of the carrier only (SSB Noise). Many sources of noise add signals on both sides of the carrier that are equal in amplitude and it is usually safe to assume the spectrum is symmetric. This does not have to be the case, however. The noise above the carrier frequency could be different to that below the carrier – the spectrum can be asymmetric.

Most of the measurements methods rely on demodulating the carrier as either phase modulation or frequency modulation – which is, by normal definition, double-sided signals with a symmetric spectrum. The measurement method assumes that the spectrum is symmetric, and measures the combined signal levels. The results take into account the fact that both sidebands are being measured and that the sidebands are correlated.

1.4 Amplitude noise

Amplitude noise is not directly covered by this guide, but it is important that the reader understands its impact. Phase (or frequency) noise is a noise signal that changes the instantaneous phase of the carrier away from its theoretical (ideal) value. It does not change the amplitude of the carrier. Amplitude noise affects the amplitude component of the signal only – it does not effect the phase of the signal.

Some noise sources add both amplitude and phase noise, the prime example being white additive noise (e.g. from a buffer amplifier). White noise added in this way is resolved into separate components when it is measured in a system – an amplitude component and a phase component, each at -3 dB to the original white noise level.

In many systems, amplitude noise is either equal to or (usually) less than phase noise for good technical reasons. Phase noise tends to increase at 6 dB per octave of reducing offset frequency from the carrier for oscillators limited by thermal noise amplified by the resonant circuit; amplitude noise does not tend to behave in the same way.

There are some systems where amplitude noise may be worse than phase noise, and this can cause some confusing results depending on the method of measurement. A classic example of

amplitude noise being higher, is where a system has a wide bandwidth amplitude leveling system that adds amplitude noise to the carrier but may not add phase noise³.

1.5 Comparing phase noise profiles

The phase noise performance of sources can differ considerably according to the type of oscillator used and its application. This variation in performance leads to very different demands on a test system designed to measure the noise – measuring a high performance crystal oscillator is very different to measuring a voltage controlled oscillator (VCO).

The following plots show just how much variation in performance can be expected.

The plot below is for a high performance crystal oscillator⁴ that is designed to have a low aging rate. The crystal is run at relatively low powers to preserve the long-term stability, leading to very good close to carrier phase noise, but more modest performance at high frequency offsets. Other crystal oscillators might be run at higher powers to reduce higher offset noise, or even include additional crystal filtering to reduce high frequency offset noise.

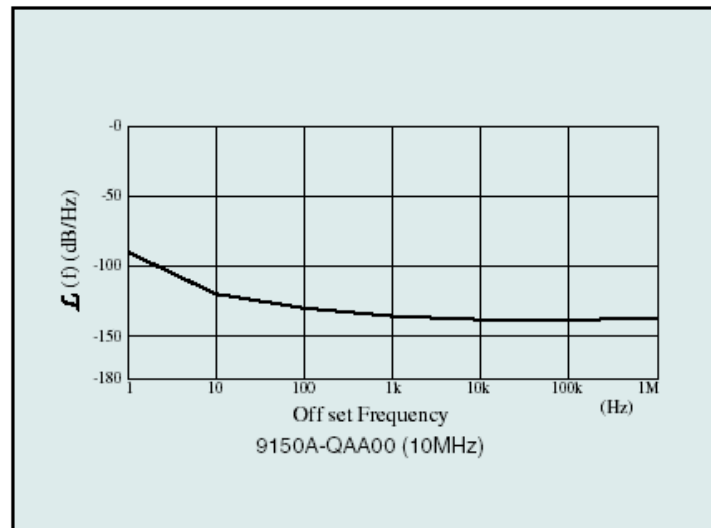


Figure 3: Phase noise of a high performance crystal oscillator

The noise performance is a smooth profile (a stylized plot) with very few features (changes of slope of the phase noise are gradual). Even at 100 Hz offset the phase noise of this oscillator is -130 dBc/Hz.

The plot below shows the phase noise profile for a relatively low cost signal generator⁵ at 1 GHz which has a very different profile. The phase noise levels are generally higher (as would be expected from a broadband general purpose source), and have more structure. For offset frequencies of 1 MHz the phase noise is dominated by broadband floor noise from the signal generator's relatively complex output systems. At lower offsets, the phase noise rises at 6 dB per octave, typical where the noise from the oscillator is due to white FM noise. Since this example is a synthesized signal generator based on fractional N technology, at approximately

³ If the circuit that controls the level does not cause unintentional phase changes this will be the case. If the circuit does it may be a source of phase noise.

⁴ Source. NDK data sheets for 9100 series

⁵ Source. Aeroflex data sheet for 2025 signal generator

3 kHz offset the phase noise reaches a plateau because the PLL bandwidth starts to remove phase noise from the oscillator. There is even a hint that the phase noise starts to lower again as the PLL gain rises at 12 dB/octave through the use of a Type 2 phase locked loop. At lower offsets the phase noise starts to rise again as noise from the frequency standard (and the dividers and other circuits in the PLL) is multiplied up, giving rise to another low frequency offset source of phase noise.




Figure 4: Phase noise of a low cost signal generator

The profile is quite different from a low noise signal generator's (costing considerably more) shown previously, where combinations of broadband and narrow-band PLLs generate a very different noise profile, reflecting the very different architectures they use.

PLL based sources can also have relatively simple profiles, as the plot for a simple OEM fractional N synthesizer⁶ below shows.

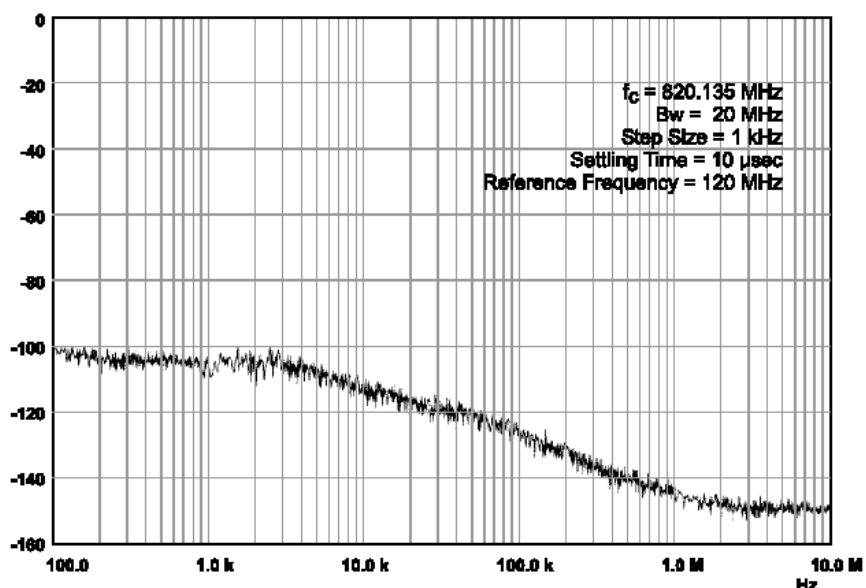


Figure 5: Phase noise of a simple OEM fractional N synthesizer

⁶ Source. Synergy Microwave Corporation

The phase noise profile is relatively simple, but the source is not expected to provide frequency modulation capability, so the PLL architecture is considerably simpler than that for a signal generator.

Broadband Voltage Controlled Oscillators (VCO's) used for general purpose applications that are available commercially and are designed for low cost, are likely to have significantly worse phase noise than the plots shown above. An oscillator operating at about 500 MHz and having a frequency cover of 25 MHz designed for 5 V only operation, could have the phase noise performance shown in the table below

| | | | | |
|------------------|------------|------------|-------------|-------------|
| Offset frequency | 1 kHz | 10 kHz | 100 kHz | 1 MHz |
| Phase noise | -75 dBc/Hz | -98 dBc/Hz | -118 dBc/Hz | -138 dBc/Hz |

The tuning sensitivity of such an oscillator on the varactor diode has to be very high, perhaps 20 or 30 MHz/volt. This makes it inevitably a relatively noisy source, but a convenient and cost effective solution for modern applications where (as in this case) 5 V operation may be important. A simple (and very rough) calculation shows that it would be hard to expect very much better performance given the tuning sensitivity:

The noise of a 50 ohm resistor at ambient temperature is very approximately 4.5×10^{-10} V in a 1 Hz bandwidth

A tuning sensitivity of 30 MHz/V generates 135×10^{-4} Hz deviation in a 1 Hz BW.

From the narrow band modulation theory 1 Hz deviation at 1 MHz offset is -126 dBc/Hz

Scaling the numbers 135×10^{-4} Hz deviation could produce approximately -163 dBc/Hz at 1 MHz offset.

At 10 kHz offset this would produce a limit -123 dBc/Hz from this one effect alone.

Although there may be clearly room for improvement, in practice other considerations will limit the performance (the Q of the oscillator, higher varactor source impedance).

There is a great deal of information available on the web about the theory of oscillator design and a more theoretical approach to phase noise. A useful starting point is the Synergy Microwave Corporation web site: <http://www.synergymwave.com/Articles/Articles.htm>, a company founded by Ulrich Rohde (also a co-founder of Rohde and Schwarz). The Agilent web site www.agilent.com has useful application notes on phase noise measurement and oscillators. The Aeroflex web site www.aeroflex.com has useful information on their phase noise measurement systems.

Measuring such widely divergent oscillators causes considerable measurement problems and it is not surprising that no one technique solves all the problems. Different test methodologies may be more appropriate to some applications than others, as will be shown in the later chapters of this guide.

1.6 Relevance of phase noise to communication systems

Phase noise is of particular interest to communication systems and the following section provides a very broad (and brief) description of its impact.

1.6.1 Analog communication systems

For analog audio communication systems the most important offset frequencies are those around 1 kHz since this strongly influences the residual FM, and therefore the ultimate signal to noise ratio⁷, of the transmitted or received signal. The noise can be heard as an audio “hiss” in the audio bandwidth of the systems. The phase noise, whether generated in the transmission systems or the receiver electronics and local oscillators, can be measured as a residual modulation (phase modulation or frequency modulation) in a specified bandwidth.

At the higher frequency offsets phase noise affects transmitted power into other RF channels, while the local oscillator noise in the receiver may affect the amount of power unintentionally mixed back into the receiver bandwidth from other sources.

1.6.2 Digital communications

Phase noise characteristics are important for digital, as well as analog, communication systems. The 1 kHz phase noise characteristics of oscillators in transmitters using Time Domain Multiple Access (TDMA) or Time Domain Duplex (TDD) techniques often strongly influences the residual phase or frequency jitter within a single burst of the carrier frequency. As wider bandwidth systems are adopted, phase noise at larger offsets will become increasingly specified, but in general, the toughest target is likely to remain the 100 Hz to 3 kHz offset performance.

The sensitivity to the noise in the 1 kHz offset region in digital modulation systems arises because the signal is split into blocks of information, typically with a duration of 1 ms to 20 ms, for the purpose of encoding speech or adding error correction. The details of this are beyond the scope of this guide. The blocks of information usually have within them, a sequence of digital bits that are used to extrapolate the phase and frequency reference of the transmitted signal over the entire block. Having obtained this phase reference, the digital data can be derived. This phase or frequency estimation process means that phase noise at low carrier frequency offsets is removed, whereas noise at frequencies corresponding to times shorter than the data block length can directly lead to an increase in measured modulation error. The longer the length of the data block used, the more susceptible the system is to lower frequency noise.

The varying amplitude of many digital communication systems can lead to the generation of noise like signals; in fact spread spectrum systems are designed to look like noise when viewed on a spectrum analyzer. Any non-linear behavior in the source with this sort of modulation will lead to spectral spreading of the signal. This type of signal should not be confused with phase noise.

Irrespective of whether a communication system is digital or analog, the measurement of phase noise is likely to continue to be an important activity in the design of communication systems, whether they are analog or digital.

⁷ More information on signal to noise and distortion ratio can be found in an application note “All you need to know about SINAD measurements using the 2023” by David Owen which is available on the Aeroflex web site www.aeroflex.com

2 General principles of methods of measuring phase noise

This guide has chapters on five basic measurement techniques based on

1. Spectrum Analyzers,
2. Delay Line Discriminators
3. Quadrature Technique
4. FM Discriminators.
5. Direct Digital Measurement

All of these methods can be used to successfully measure the characteristics of a signal source and they each have their advantages and disadvantages. There is overlap between all these methods, but as far as possible, they have been kept separate to give a more structured approach.

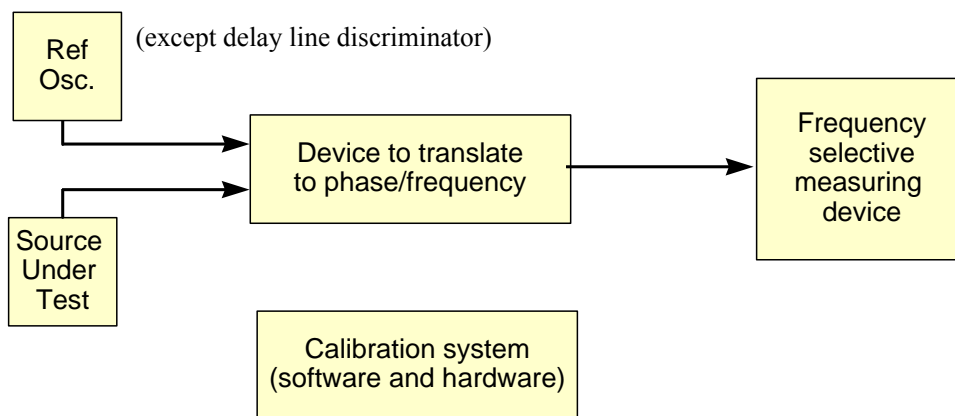


Figure 6: Components of a phase noise measurement system

The first four methods of measurement all rely on a similar basic principle. The signal to be measured is frequency converted to a baseband or IF and then passed through a device which extracts either phase or frequency information from the carrier. A frequency selective measuring device is then used to measure the noise as a function of offset frequency. A calibration system is used to scale the results into meaningful units.

In the case of measurements with spectrum analyzers, the local oscillator can be the local oscillator of the spectrum analyzer.

In the case of the delay line discriminator method, the source under test is used for as its own local oscillator.

The frequency selective measuring device can be a spectrum analyzer, an FFT analyzer or a high speed digitizer whose output can be converted into a power (or voltage) versus frequency display. In this guide, the term spectrum analyzer has been used to describe both spectrum analyzers and FFT analyzers.

The chapter on the use of digitizers has been inserted to describe the use of fast digital acquisition systems that could be considered to include FFT analyzers.

3 Spectrum analyzer methods

This section describes using an RF spectrum analyzer to measure phase noise; it is not intended to cover the use of FFT analyzers though in principle they could be used in some setups.

Since Spectrum Analyzers measure the RF signal power in a specific bandwidth, they can clearly be used to measure phase noise. Most modern analyzers include software functions which will convert a measured signal level from its measured value (in the spectrum analyzer filter bandwidth) to the equivalent noise signal in a 1 Hz bandwidth provided the noise can be treated as Gaussian and flat within the bandwidth of the filter.

The measurement of phase noise can then be simply be a question of connecting the unit under test to a spectrum analyzer and making the required measurements




Figure 7: Phase noise measurement with a spectrum analyzer

By measuring the total carrier power (if necessary on a wide filter setting) and then measuring the noise signal at a specified offset from the carrier, a (normalized to a 1 Hz bandwidth) phase noise measurement can be derived. Using a wider bandwidth filter to measure the carrier power can avoid problems with close to carrier noise or frequency drift in the UUT or in the local oscillators of the spectrum analyzer from causing errors in finding the true peak of the signal.

Many spectrum analyzers provide functions specifically designed to help perform this measurement, in particular:

- Peak find measurements
- Peak tracking
- Normalization to 1 Hz bandwidth
- Dual markers with relative function
- Use of marker with a “relative” function

It is also important that a user of spectrum analyzer undertaking a phase noise measurement understands that the system does not just measure phase noise. A spectrum analyzer plots the total signal power seen by its filter as a function of frequency. This signal contains both amplitude and phase noise, and the sum of the two is displayed on the screen. The spectrum analyzer cannot distinguish between the two. There are cases where amplitude noise is higher than phase noise components, but it is not the most common situation and is usually true only over limited offset frequencies.

Using a spectrum analyzer allows the user to see asymmetric phase noise spectrum in a way that is not easy with alternative measurements methods.

In practice, the performance of simple spectrum analyzer measurements is limited for reasons that become apparent when the block diagram of a typical spectrum analyzer is examined. However, a spectrum analyzer is a relatively common item of test equipment in many laboratories, and if one is available, and the measurement required to be made is not too demanding, a spectrum analyzer may be the natural choice.

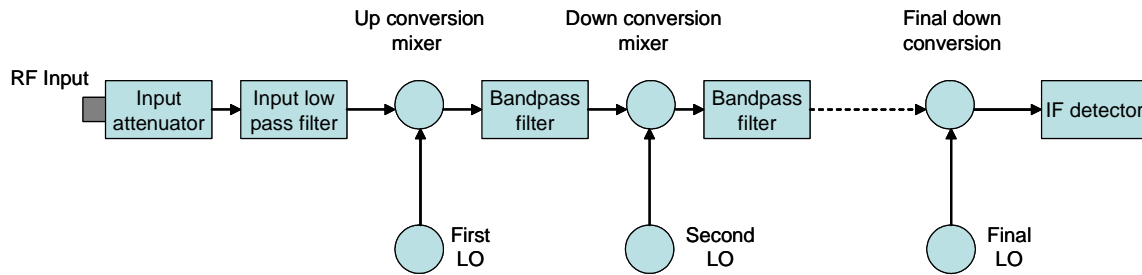


Figure 8: Block diagram of a typical spectrum analyzer

The input signal to the spectrum analyzer is first conditioned by filters and an attenuator before being up-converted to a higher frequency by a local oscillator. The resulting signal is then filtered and down-converted using a series of local oscillators and mixers to a lower frequency, where the signal can be measured. The signal is measured after being passed through one or more band pass filters. The IF detector will include a selection of band pass filters that determine the spectrum analyzer's measurement bandwidth. In some cases, the filters are constructed from inductors, capacitors, crystals or ceramic devices. However, increasingly these filters are being replaced by digitally derived filters applied after the low frequency signal has been converted from a digital to an analog signal.

Depending on the filter BW being used, the final IF used can be changed to make the design of the required filter practical (larger bandwidths use higher IF's).

Note: At higher frequencies the spectrum analyzer typically uses a harmonic sampling system and pre-selectors to reject image signals. This will make the noise levels for the spectrum analyzer local oscillators higher, often producing a “stepped” change of performance with frequency.

To provide the swept frequency measurement, one or more of the local oscillators needs to be swept in frequency so that the measurements of power are made at a fixed lower frequency. The local oscillator must be designed to be frequency agile to meet the main objectives of the product – and that limits the performance of the spectrum analyzer's local oscillator in terms of phase noise.

A spectrum analyzer cannot be used to measure phase noise levels below the inherent performance of the combined local oscillators in the instrument without additional hardware.

There are also other limitations in a spectrum analyzer. The product is designed to produce a swept frequency plot of a signal over a very wide frequency range without the introduction of artefacts that are not present on the real signal. To ensure good performance the input signal

levels to the mixers need to be carefully controlled to avoid intermodulation products. This tends to limit the wide band signal to noise ratio of the spectrum analyzer, in turn limiting the noise measurements that can be performed far from the carrier.




Figure 9: A phase noise measurement on a spectrum analyzer

The spectrum analyzer filters further limit the capability of the measurement. The carrier is displayed as a large signal and on either side the filter response can be seen. It is not until the user gets to a considerable offset that the ultimate performance of the analyzer can be seen. This limits the capability of the spectrum analyzer to measure phase noise close to the carrier frequency.

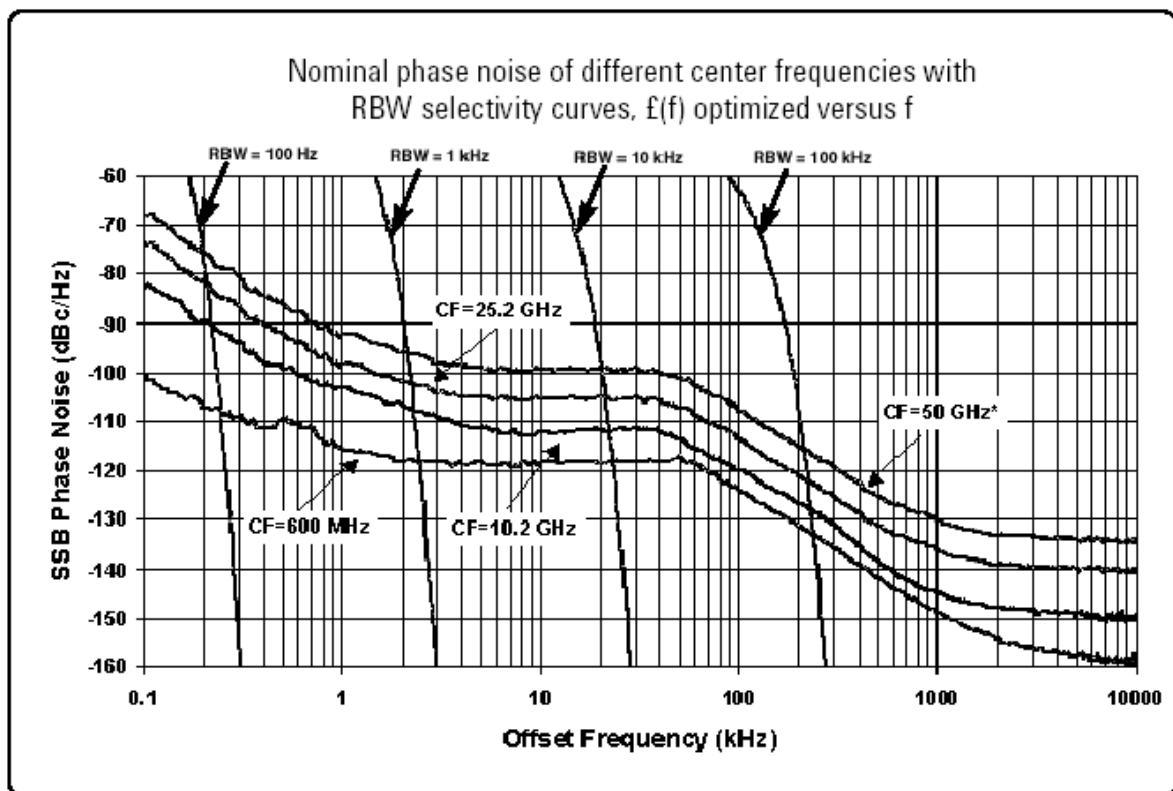


Figure 10: Limitations of phase noise measurements with high performance spectrum analyzer⁸

The net effect of these limitations can be seen graphically on the above diagram for a high performance spectrum analyzer (lower performance spectrum analyzers will place even more limitations on the measurement capability).

For frequencies close to the carrier, the measurement is limited by the spectrum analyzer filter response. Selecting the 100 Hz filter limits measurements to broadly those above 300 Hz offset. A narrower filter will improve this minimum offset frequency on a roughly pro-rota basis.

At greater offsets, the performance is limited by the phase noise of the local oscillator system. These limitations will be frequency dependent, the higher the frequency the greater the limitations.

Note that in this example the performance at 600 MHz is reasonably representative of a wide range of input frequencies, including relatively low frequencies. Until the spectrum analyzer starts to use harmonic input sampling, the performance is relatively frequency independent.

At higher frequency offsets the phase noise limitation is imposed by the signal levels permitted at the mixers and the noise limit flattens out. As the input frequency is raised, the harmonic sampling further limits the spectrum analyzer performance.

The performance of the system at the larger offsets can be improved by ensuring the input attenuator to the spectrum analyzer is correctly set, and the setting is optimized for dynamic range in preference to linearity.

Those spectrum analyzers which perform narrow band measurements using digital techniques, can generally perform better close to carrier measurements than analogue versions and, with the right software, provide more reliable conversion of measurement results to phase noise.

3.1 Down converter and IF filter with a spectrum analyzer

The major problem that a spectrum analyzer has in measuring phase noise is that, unlike some of the other methods, the carrier has not been removed from the signal. The spectrum analyzer is forced to make a high dynamic range measurement that exposes limitations in the performance.

For some applications the noise limitations of spectrum analyzers can be overcome by the use of band-pass filters or band-stop filters.

The method requires the use of a second reference RF or microwave signal source and a mixer to convert the signal to a convenient intermediate frequency (IF). The signal from the mixer is then passed through a low pass filter to reject the sum components from the mixer, a low noise amplifier to increase the signal level and a band-pass crystal (Xtal) or ceramic filter

⁸ Source. Agilent data sheet for PSA range of spectrum analyzers

before being measured by the spectrum analyzer. The band-pass filter can be a commercially available inductor/capacitor, crystal or ceramic IF filter of the type commonly used in receiver systems. The filter needs to be able to reject the IF frequency but pass a selected band of noise to the spectrum analyzer. The filter needs to have a flat pass band response to minimize the calibration problems that can be introduced. The spectrum analyzer filter bandwidth is set to be narrower than the band-pass filter and it is this filter that determines the bandwidth of the measurement.

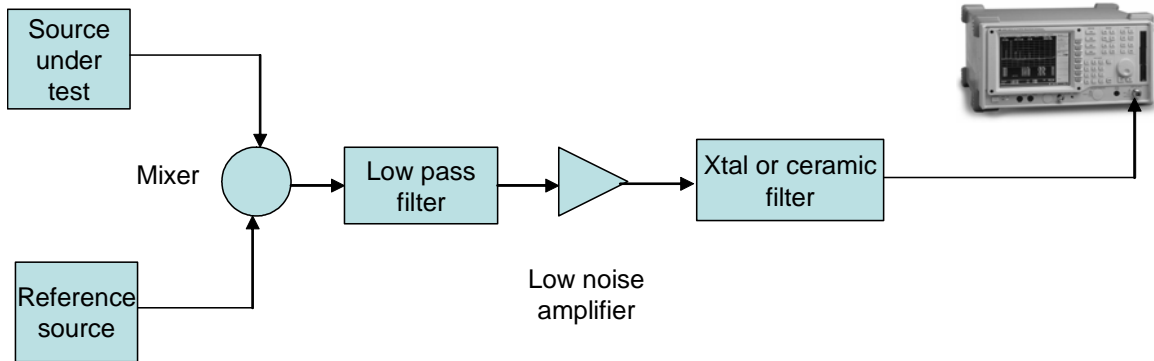


Figure 11: Phase noise measurement using a down converter and a spectrum analyzer

The improvement in performance using this method occurs because the carrier frequency is rejected at the IF output by the band-pass filter. The spectrum analyzer is being used to measure a signal that is almost entirely consists of the signal that is being measured (the phase noise). Noise introduced by the mixing of the signal from the unit under test with the spectrum analyzer local oscillator is substantially reduced, as shown below, so the restrictions on measurement dynamic range are no longer critically dependent on the quality of the spectrum analyzer’s local oscillator. The dynamic range required of the spectrum analyzer is significantly reduced.

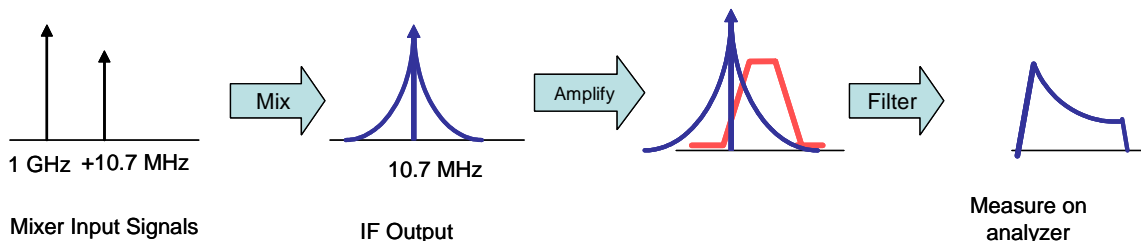


Figure 12: Frequency diagram for phase noise measurement using a down converter and a spectrum analyzer

Some care needs to be taken in making measurements in this way. Suitable filters with narrow bandwidths are rarely designed for 50 ohm systems and often have severe changes of impedance with frequency. The mixer has to be buffered from this impedance variation to avoid errors due to reflected signals re-mixing. The filters can also exhibit non-linear behavior at both low levels (particularly crystals) and high levels (as crystal or ceramic devices exceed their linear power ratings). These problems, combined with frequency response unflatness in the pass band, can make the measurement accuracy unreliable unless basic precautions are taken.

The technique is also restricted to measurements at offsets of typically greater than 10 kHz since it relies on the filter having to reject a significant proportion of the carrier signal at the IF.

A variation on this approach is to use a band stop filter to reject the IF signal, but such filters are not as commonly commercially available. To use a band stop filter, the IF is centered in the middle of the band-stop characteristics to suppress the carrier level while allowing the spectrum analyzer to measure the noise.

3.2 Spectrum analyzer measurement calibration

A great advantage of the spectrum analyzer method is that, if it is an appropriate method, in general the manufacturer's data sheet provides a great deal of guidance on calibration. However, users should be aware of the limitations on spectrum analyzer accuracy described in the section on estimating uncertainty.

When a spectrum analyzer is used to measure phase noise using a filter to improve dynamic range, the calibration process is slightly more complicated, but not onerous. The complication arises because, unlike the direct measurement, the carrier has been partially removed, so there is no reference.

For systems using a band-pass filter:

- Set the difference frequency between the UUT and the LO to be equal to the centre of the band pass filter response.

- Measure the level of the IF signal, applying any input attenuation to the spectrum analyzer that may be required to prevent it from overloading

- Offset the LO frequency by the offset frequency that the phase noise is to be measured at.

- Measure the phase noise level relative to the carrier using the same spectrum analyzer resolution filter settings to avoid additional errors in accuracy.

Offsetting the LO by the offset frequency of the phase noise measurements will reduce errors caused by the band pass filters frequency response (ripple) provided the filter has no significant non-linearity. Keeping the spectrum analyzer's resolution filter the same will minimize errors introduced by the switching of resolution filters.

The dependence of the system on the amplitude of the source under test can be reduced by making the LO input to the mixer the source under test rather than the reference source provided there enough power to drive the LO port satisfactorily. The level variation that can be accepted is rather limited by the need to drive the mixer over a relatively limited amplitude range for correct operation. Using this method will also remove some amplitude noise.

4 Delay line discriminator

The delay line discriminator method relies on demodulating the signal from the UUT to provide an output which is representative of the FM noise of the source.




Figure 13: Block diagram of a delay line discriminator phase noise measurement system

An FM detector is constructed by taking the RF signal to be measured and splitting it into two paths. One path is passed directly into a mixer and the second path is passed through a delay line. The two outputs are mixed together in a double balanced mixer, the sum frequency being removed by a low pass filter. The mixer has identical frequencies applied to its LO and RF ports so the output from the mixer is essentially a DC voltage whose level is dependent on the phase difference between the two signals.

The delay line, the bulk of which is usually implemented as a low loss coaxial cable, includes a variable phase shifter or a mechanically adjustable transmission line so that the phase of the two signals applied to the mixer can be set for quadrature. When the two signals are 90 degrees out of phase the mixer output is nominally 0 Volts.

The measurement system behaves as an FM discriminator, a fact that can be understood without the use of mathematical formula. If you consider a source which has a single sine wave phase modulation applied to it, the system delays the signal of one path compared to the other. The two FM signals will have a different phase at the input to the mixer, and this will be reflected by a sine wave voltage appearing at the output of the mixer whose amplitude is dependent on the phase difference between them. As the frequency of the modulation is reduced, the level of the demodulated signal becomes less, because the time delay becomes a smaller proportion of the sine wave phase modulation (in other words the signal level changes less in a given time interval for lower frequencies). This is precisely the behavior you would expect of an FM discriminator.

The input ports to the mixer can be driven at high RF levels to make the output response a linear phase detector rather than a conventional sinusoidal output. The mixer could alternatively be driven with a lower level to the RF port, but the reduced output from the mixer makes the system less sensitive.

In practical systems, at least one amplifier is likely to be needed to make sure the mixer is driven adequately.

A full analysis leads to the output being of the form:

$$V_{\text{noise at } f_m} = K 2\pi T (\Delta f_m) (\sin x)/x$$

K is the sensitivity of the phase detector, expressed in Volt/radian, which in turn is dependent on the mixer drive and loss

f_m is the offset frequency

Δf_m is the FM deviation (Hz) at the offset frequency f_m

T is the delay line delay

and x is $\pi T f_m$

The bandwidth of the discriminator is a classic Sine(x)/x response with the first null at a frequency equivalent to the time delay between the two RF paths. The null frequency occurs because, at this frequency, the delay represents a whole cycle of any phase modulation that might be present.

The sensitivity of the FM discriminator formed by the system is dependent upon the RF level applied, the conversion loss of the mixer and the time delay of the delay line. The longer the delay line the greater the sensitivity of the measurement but the more restricted its measurement bandwidth since the first null of the Sine(x)/x response is lower.

A significant advantage of this technique is that it does not require the use of a second RF source to convert the frequency of the source to be measured to a fixed IF (or base band signal). This removes one potential source of error, i.e. an additional source of noise from the reference oscillator. Also, since the method is based on the use of a frequency discriminator it is not very prone to being overloaded by low frequency sources of phase noise (e.g., power supply related signals). If the frequency of the UUT varies with time, it does not significantly impact the operation of the system, since it simply generates a DC offset at the output. Provided this is small relative to the bandwidth of the discriminator, it will have little effect on the system performance.

In order to drive the mixer efficiently the system often needs amplifiers to manage the two RF paths. At least one of the paths needs a signal large enough to drive the mixer adequately so that there is not an excessive conversion loss.

The delay line is typically implemented as a coil of coaxial cable. As the frequency of the source is increased, the attenuation of the coaxial cable increases, and it is more likely that additional amplifiers will be needed to maintain the sensitivity of the system.

The mixer behaves as a phase detector, but the system measures FM noise – the output from the system is dependent on Hz deviation (in each 1 Hz of measurement bandwidth in the case of noise signals). This figure has to be converted to a phase noise measurement rather than frequency noise measurement using the methods described previously.

4.1 Calibration of delay line discriminators

The delay line discriminator needs to be calibrated for each measurement. The calibration is dependent on a number of factors, including the RF levels, the insertion loss of the two paths and the length of the delay line. These factors change between sources with time, and with other changes in the setup, so calibration has to be performed each time an accurate measurement is made. Calibration of the system is not entirely straightforward.

A simple method of calibration is to use a substitution technique – the source under test is replaced with a signal generator that can produce calibrated FM (or phase modulation). The signal generator must have the both the same frequency and, more importantly, the same RF level as the UUT. To calibrate the system a known level of sine wave FM is set on the signal generator and the output from the system is measured as a sine wave (at the modulation frequency). Since the FM signal is known (and the dBc on the carrier can be calculated from narrow band modulation theory) the user has a calibration figure for the entire system. Accuracy is then primarily determined by the accuracy of the substitution, the accuracy of the FM and the accuracy of the spectrum analyzer used to derive the frequency versus noise information. More information on the use of a signal generator used this way is provided in the section on measurement uncertainty.

The more commonly used method of calibration is to derive a calibration factor for the amplitude characteristics of the system and adjust it for the time delay between the two paths. The amplitude information is obtained by adjusting the time delay of the second path (by a sliding transmission line or electronic phase shifter) until the mixer DC output level finds the peak positive and negative voltage. From this, the phase detector sensitivity can be deduced provided the waveform shape (triangular or sine wave) is known.

Using the formula:

$$V_{\text{noise at } f_m} = K 2\pi T(\Delta f_m)(\sin x)/x$$

referred to above, the remaining unknown is the delay line length which has to be independently measured or calculated. With V_{noise} being measured, the deviation at a specified offset can be calculated (Δf_m) in Hz and converted to phase noise.

If the system is measuring phase noise at large offsets the correction for the $\sin(x)/x$ curve must be applied, otherwise the results will look optimistic. For convenience, a table of correction factors is given below.

Table of correction values versus offset frequency, where offset frequency is expressed relative to the first null frequency ($1/T$)

| Offset frequency/first null frequency | Sinx/x response | Correction that needs to be applied (dB) |
|--|----------------------------|---|
| 0 | 1 | 0 |
| 0.02 | 0.999 | -0.006 |
| 0.04 | 0.997 | -0.023 |
| 0.06 | 0.994 | -0.051 |
| 0.08 | 0.990 | -0.092 |
| 0.1 | 0.984 | -0.143 |
| 0.12 | 0.976 | -0.207 |
| 0.14 | 0.968 | -0.282 |
| 0.16 | 0.958 | -0.369 |
| 0.18 | 0.948 | -0.468 |
| 0.2 | 0.935 | -0.579 |
| 0.22 | 0.922 | -0.703 |
| 0.24 | 0.908 | -0.839 |
| 0.26 | 0.892 | -0.988 |
| 0.28 | 0.876 | -1.151 |
| 0.3 | 0.858 | -1.326 |
| 0.32 | 0.840 | -1.516 |
| 0.34 | 0.820 | -1.719 |
| 0.36 | 0.800 | -1.938 |
| 0.38 | 0.779 | -2.171 |
| 0.4 | 0.757 | -2.420 |
| 0.42 | 0.734 | -2.685 |
| 0.44 | 0.711 | -2.967 |
| 0.46 | 0.687 | -3.267 |
| 0.48 | 0.662 | -3.585 |
| 0.5 | 0.637 | -3.922 |

5 Quadrature method

In the Quadrature System, two oscillators at identical frequencies are used.

One of the oscillators will be the source being tested and the other will be a reference source whose performance is known to be better than the source under test. The sources are combined in a mixer and the resulting output signal is filtered and amplified by a Low Noise Amplifier (LNA). A Fast Fourier Transform (FFT) Analyzer or a Spectrum Analyzer typically measures the output from the mixer.

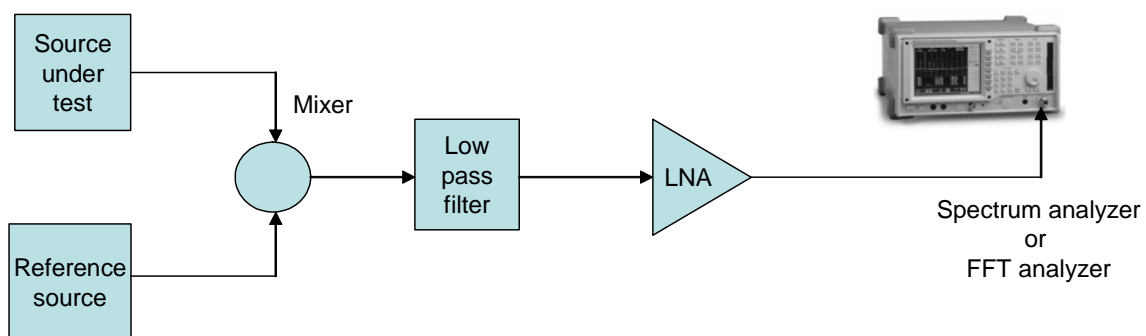


Figure 14: Block diagram of a phase noise measurement system using the Quadrature method

The system behaves as a phase detector if the input signals to the mixer are adjusted to be in phase quadrature. When they are in quadrature the nominal DC output voltage from the mixer is 0 Volts. The output from the mixer is a direct measure of the phase difference between the two signals, and so it is also a direct measure of the phase noise of the UUT (assuming the reference source is better than the UUT).

Setting the sources to be in phase quadrature is not always very easy. If both frequency sources are synthesizers with good long-term stability, then there is usually not a great problem. Phase quadrature can usually be achieved by offsetting the frequency of one of the sources, during which time the system will produce an AC voltage at the same frequency as the offset frequency, and then resetting the frequency to be the same. The synthesizers will settle to a new phase that can be checked to see if it is close enough to phase quadrature, typically a voltage less than 5% of the peak voltage from the phase detector when the frequencies are offset⁹. If one of the sources has fine frequency control (e.g. 0.1 Hz) the frequency can be set to a small offset and then returned to the required frequency when the correct phase is achieved. Alternatively, many modern synthesized sources include phase adjustment controls that can be used to adjust the phase more easily.

However, in the many applications, measurements are undertaken under less than ideal conditions, and a feedback system has to be used to maintain phase quadrature. The feedback system forms a phase locked loop that controls the frequency/phase of one of the oscillators to

⁹ This assumes the signal is larger than any residual offsets from the phase detector. More on this can be found in the section on measurement uncertainty.

correct for departures from quadrature. Essentially the phase locked loop acts to drive the output voltage of the mixer to 0 Volts. Often a Type 2 PLL is used, as shown, so that the gain of the loop at DC is infinite.

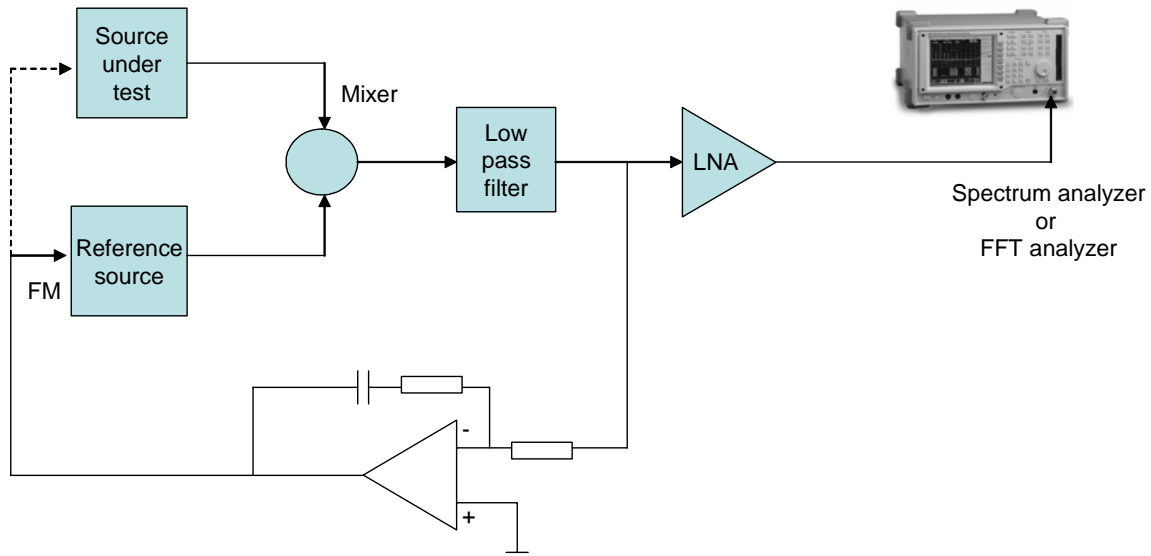


Figure 15: Block diagram of a phase noise measurement system incorporating a phase locked loop to maintain quadrature

The use of a phase locked loop to maintain phase quadrature does imply some knowledge of the tuning characteristics of one of the oscillators and the mixer drive levels, since the bandwidth of the phase locked loop is affected by both of these parameters. Knowledge is required to ensure the PLL is stable and has a bandwidth that allows the measurement to be performed. The break frequency introduced in the feedback integrator (that makes it a Type 2 Loop) needs to be low enough to ensure stability but not so low as to cause the loop to have an excessive settling time. A working knowledge of phase locked loop design is clearly an advantage.

In many applications, the availability of a low noise signal generator with a high performance DC coupled FM capability can be extremely helpful. It allows the signal generator to be used as a highly controllable VCO (or VCXO). The tuning slope of the “virtual” VCO can be increased or decreased by the simple expedient of changing the FM deviation setting, which in turn adjusts the PLL characteristics (increasing or decreasing the PLL bandwidth).

If the peak phase excursion of the noise exceeds 0.1 radians, the mixer phase detector response becomes non-linear and degrades the measurement accuracy.

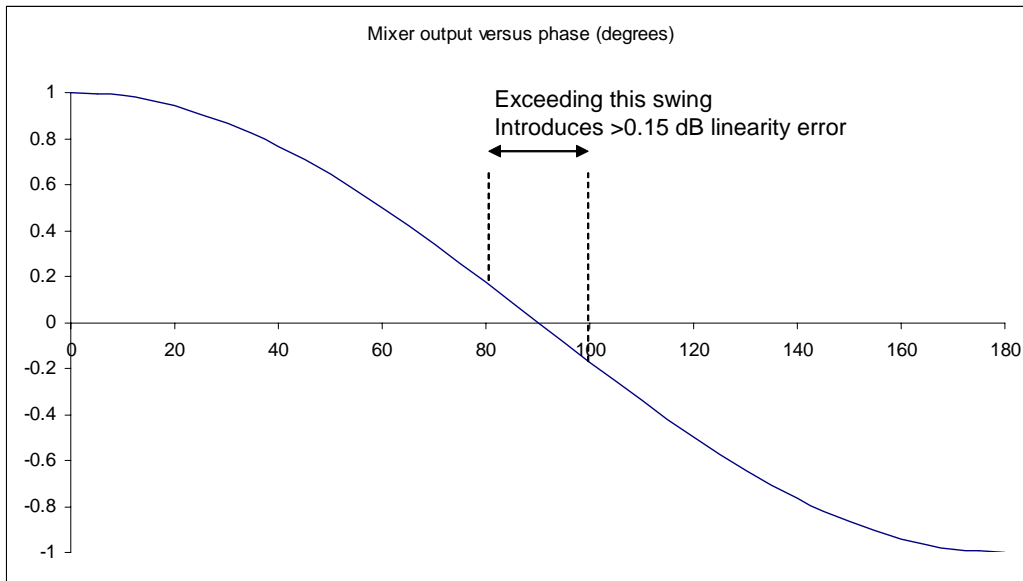


Figure 16: Mixer output versus phase

Since the peak phase excursions are caused primarily by low frequency noise then, under these conditions, the phase locked loop bandwidth can be widened in order to restrict the peak phase excursion. This will complicate the calibration of the system since ultimately the effects of the PLL must be removed from the system results.

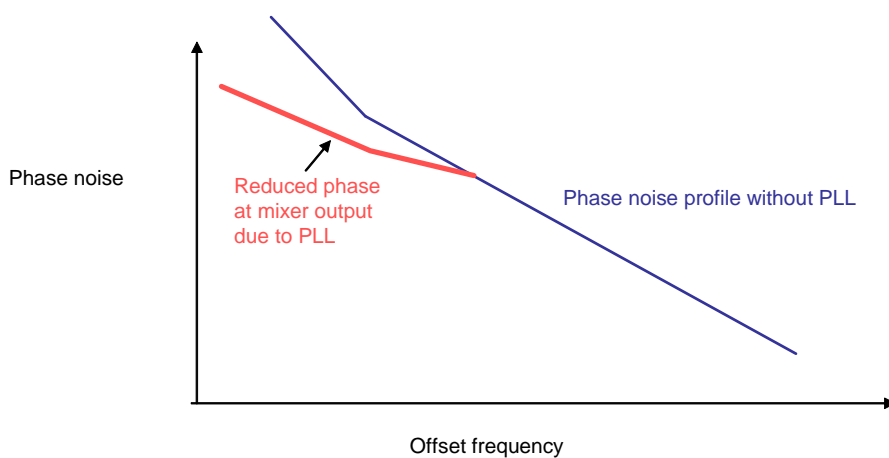


Figure 17: Phase Locked loop suppression of measured phase noise

The sensitivity of the Quadrature Method is extremely good for both low and high frequency offset noise measurements. It is often used to measure very high performance oscillators, including crystal oscillators and atomic frequency standards. For applications requiring the measurement of very low noise levels, particularly at large frequency offsets where they are more likely to occur, there may be a need to get as much performance from the test system as possible. This will generally require the drive levels to the mixer being high on both ports – driving the system components at high level will generally produce higher signal to noise ratios. Driving both ports of the mixer at high levels will cause the mixer to behave in a different way to driving it with the RF port level significantly below that of the LO level.

The Quadrature Method rejects amplitude noise since the mixer behaves as a limiter to the amplitude component on the local oscillator and, provided the signals are in quadrature, the

amplitude component on the linear (reference) input port to the mixer only has a very weak secondary effect on the output from the mixer.

One potential issue for the Quadrature Method is that it demodulates the phase noise rather than the frequency noise. Although it does make correction of the plot easier (no FM to phase conversion), on many sources it can lead to the phase detector or the following measuring equipment being overloaded with the low frequency noise typical of VCO's. In this case, the problem is making sure the dynamic range of the measuring equipment can cope with the required measurement. To overcome this the measurement system may have to change the bandwidth of the PLL depending on the offset frequency to be measured, so that for wide offset frequencies, the phase noise at low frequencies has been reduced to avoid overloading the phase detector and the measuring system. The PLL bandwidth can then be narrowed and the measuring equipment rescaled to measure the higher signal levels of the low frequency noise. This can lead to much more complex calibration, and increase the potential for "stitching errors" (see section on common measurement problems).

5.1 Calibration of the quadrature method

In order to carry out a measurement, the quadrature system has to be calibrated since the sensitivity of the measurement is dependent on the insertion loss, drive level used for the mixer and the PLL characteristics. There is an additional complication in systems where a PLL is used to maintain quadrature, since the PLL has a major impact on interpreting the test results. In these cases, two calibration steps may be required unless there is a high degree of confidence that the PLL is having no effect at the offset frequencies of interest.

5.2 Calibration of higher offsets

If the LO input level required for the mixer is substantially greater than the RF port drive, a calibration assessment can be obtained by offsetting the frequency of one of the sources by a small amount. A low frequency sine wave (equal to the frequency difference) is produced at the output of the mixer since the phase of the two signals is constantly changing. The amplitude can be measured to determine the sensitivity of the mixer. The sensitivity of the measuring system is equal to the peak output voltage of the sine wave (measured in volts) with units of Volts per radian. If the peak output voltage of the mixer is 1 Volt, the sensitivity of the system is 1 Volt per radian.

More generally, the phase noise is given by:

$$\text{Phase noise} = 20\log(\Delta V/V) - 6 \text{ dB}$$

where ΔV is the measured output voltage at the wanted offset (in 1 Hz BW)
and V is the sensitivity of the system (in volts per radian)

If both ports of the mixer are driven at a high level to give the maximum sensitivity, then the waveform from the mixer will be more like a triangular waveform than a sinusoid, and the mixer sensitivity should be more linear with errors in phase quadrature present. However, the slope of a triangular wave is more difficult to measure accurately than in the case where a sine wave is produced, and there can be some complication (see the section on measurement uncertainty). If a triangular wave is perfect, the slope is equal to the peak output voltage divided by $\pi/2$ (approximately 1.57). It is, however, not always a safe measurement, and there is of course the transitional area where it is neither a sine wave nor a triangular wave.




Figure 18: Phase detector responses

A further complication in the calibration process with high level drives on both ports can arise if the drive signals are not well matched to the source impedance. Where the ports of the mixer are being driven hard, the mixer tends to convert the signal to a square wave and reflections can cause re-mixing and slope perturbations in the output.

An alternative, and often more reliable method of calibration, is to use a signal generator as one of the sources and to set a known amount of phase or frequency modulation¹⁰. Measuring the resulting output can provide the required calibration information. The phase modulation applied has to have a modulation index of less than 0.1 radians¹¹ to avoid mixer overload, and a modulation frequency significantly in excess of the phase locked loop bandwidth used for setting up phase quadrature. A variable frequency modulation source should quickly reveal the true PLL BW effects, and calibration should be performed ideally a factor of 10 above this frequency so its effects can be ignored (it is possible to use a lower margin than this but since it is rare for the actual frequency of the calibration to be an issue, it is better to use a safe margin). Note that the section on calibration contains some advice on using signal generator FM drives to calibrate phase noise measuring systems.

If the signal generator noise performance is not good enough to perform the actual phase noise measurement it can be used as a temporary substitute for the low noise LO in order to perform the calibration. If this method is used it is best to do the substitution on the LO port of the mixer since differences in RF level are less likely to have an impact on the system calibration.

Another approach can also be used if a synthesized signal generator is used as a local oscillator. Many of these generators have phase advance or retard features that are digitally derived from the use of the dividers in a phase locked loop. Since these features are digitally derived, they can be inherently very accurate phase shifters, providing the manufacturer of the instrument has not had to use software rounding of the phase shift to display the phase shift in a simple fashion. Consulting the operating or service manual should reveal the true accuracy of the phase shifting system; it is often derived as a binary sub multiple of 360 degrees. If the signal generator has an accurate phase shifting feature, the calibration can be performed by choosing a phase shift, and noting the change in DC level at the output of the mixer. Note this

¹⁰ Further information on this given in the section on measurement uncertainty

¹¹ An explanation of where this figure is derived is contained in the section on measurement uncertainty

method cannot be used for systems having to use a PLL to maintain phase quadrature – the phase locked loop will remove the phase change at the mixer.

5.3 Calibration of PLL effects (close to carrier phase noise measurements)

If a PLL is used to maintain phase quadrature in the system, it can cause a major complication to the calibration of the system because it removes low frequency components (and ultimately the DC component) from the output of the mixer.

The errors introduced by the phase locked loop must either be set so that they are below the frequency offset of interest, or they have to be corrected for by measuring the loop characteristics and then mathematically correcting the measurement result.

There is a further practical problem that needs to be assessed. If there is a lack of isolation between the two RF sources then, as their frequencies are brought close together, there will be a tendency for them to become injection locked. If one of the oscillators is a VCO then this is almost certain to happen and it will need to be characterized if there is a risk it will affect the results (a fact that is sometimes hard to judge). Under these conditions, it is advisable to ensure that the deliberate phase locked loop bandwidth exceeds the injection locked bandwidth.

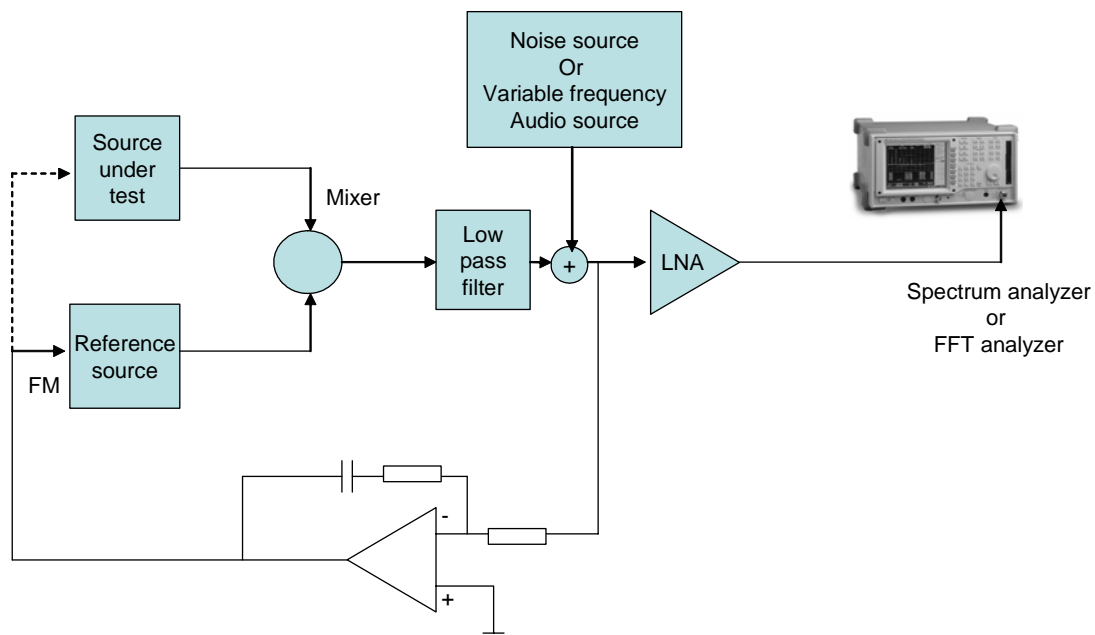


Figure 19: Block diagram of the method used to characterise the Phase Locked Loop suppression

The PLL must be characterized if accurate measurements are to be made on the source; it is rarely safe to assume that the setup does not have unintended characteristics in the phase locked loop.

The phase locked loop response can be measured by injecting a calibration signal into the loop. The calibration signal can be a swept or variable signal (e.g., the tracking generator output of a spectrum analyzer or the modulation oscillator of a signal generator) or a noise source (often available on an FFT analyzer). Outside the loop bandwidth the analyzer

measures the amplitude of the calibration signal, but inside the loop bandwidth the PLL reduces the level of calibration signal measured. From the frequency response plotted on the analyzer, a correction plot can be deduced and applied to correct the phase noise measurement results.

Even if the phase noise at just one offset frequency is required to be measured (for example against a prime specification parameter) a variable frequency source is required, since a measurement outside the loop bandwidth is required as well as a measurement at the wanted offset frequency.

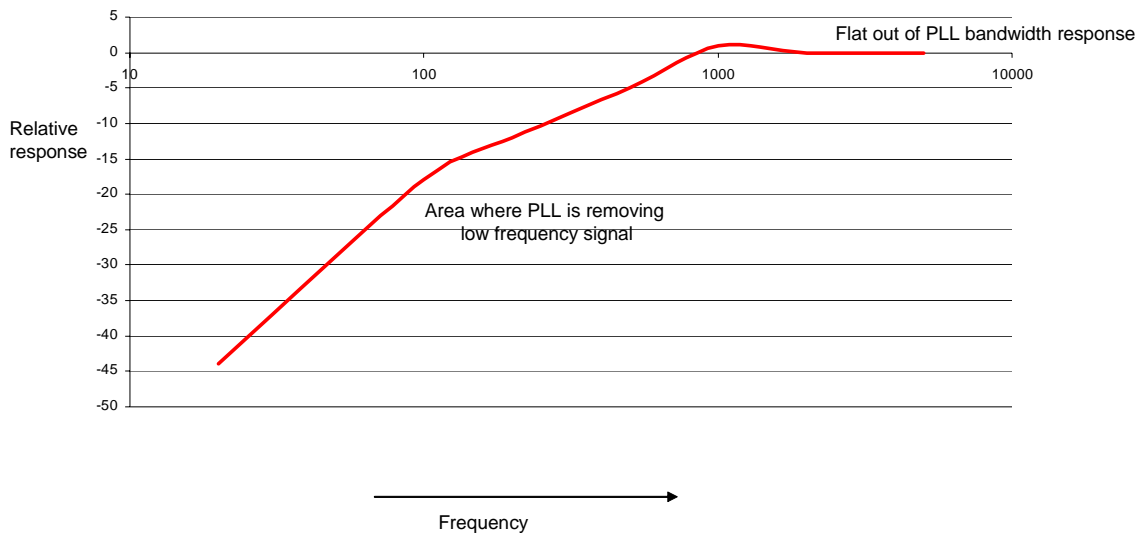


Figure 20: Typical PLL suppression characteristics

The plot of the phase locked loop response will reveal the out of band response (the flat response where the PLL is having no response) and the PLL effects at lower frequency offsets. At the edge of the PLL bandwidth the loop may exhibit a reasonably flat transitional response if the phase margins of the PLL are high, or it may have a peaked response that reverses the sign of the correction value over a range of frequencies. In general, it is safer to work with a PLL with good phase margins; the correction figures can become unduly level sensitive at the edge of the PLL bandwidth.

Care needs to be taken when interpreting results that include high correction factors (frequencies well inside the PLL bandwidth), the software may display the answers to a high degree of precision not reflected in the real accuracy of the numbers. In addition to having high correction values, there may be other effects present, such as inherent limitations in the PLL noise performance or higher risks of very variable injection locking characteristics. If very close to carrier phase noise is of interest, it is rare for the user to need large PLL bandwidths, and it may be an indication that the user has chosen an excessive bandwidth. In general it is best to use the minimum PLL bandwidth required to maintain phase quadrature over the measurement time, and to restrict the output voltage swing to less than 0.1 radian.

Some applications may use more than one setting for the PLL bandwidth in order to control the low frequency noise amplitude. Clearly, this will add some complications to the calibration of the system and the application of the correction values.

6 FM discriminator method

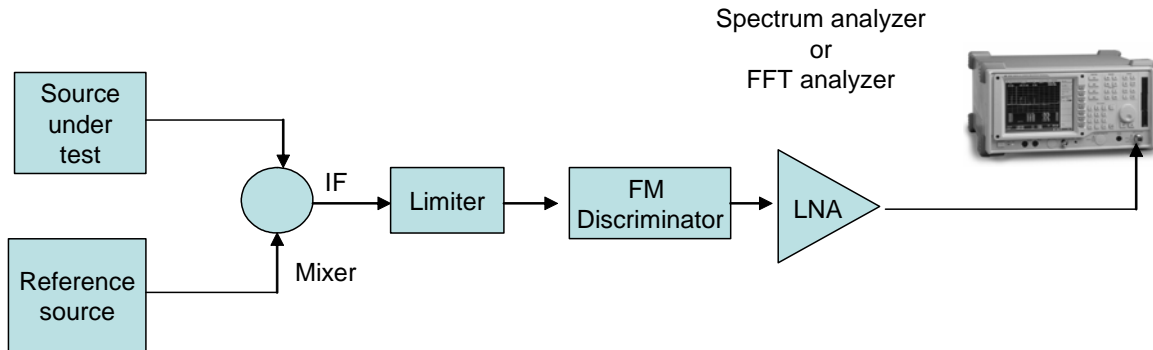


Figure 21: Block diagram of a phase noise measurement system using an FM discriminator

This method uses a mixer and a reference source to convert the signal to an IF where it is demodulated by an FM discriminator. As far as the author is aware, there are no commercial solutions available for this measurement method despite the fact that for some applications the method has some unique advantages over the alternative approaches.

In principle, any FM discriminator, including discriminators of the type found in a modulation analyzers or even an FM receiver, can be used. However, the noise performance of the discriminator is likely to have a critical effect on the ability to make a phase noise measurement and if high performance measurements are required, a carefully designed discriminator is required.

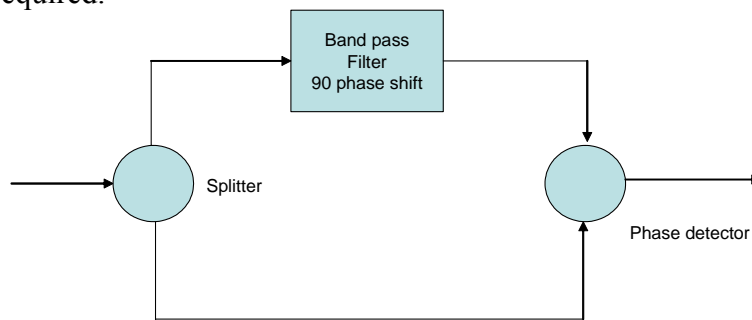


Figure 22: One suitable form of an FM discriminator

In the past the author has used a 1.5 MHz discriminator based on the use of a splitter, a band pass filter and a mixer acting as a phase detector. The band pass filter uses a coupled resonator design that ensures that at the centre frequency of operation, the phase shift through the filter is 90° , so the inputs to the phase detector are in quadrature. In the practical implementation, two band pass filters were available, one allowing a measurement bandwidth of up to 20 kHz and the other allowing measurements to 100 kHz offset (the narrow band version gave more sensitive results). Operation at an IF also allowed the FM discriminator to be implemented using a different type of phase detector operating at much higher signal levels. The design used two transformer coupled full wave rectifiers on each of the signal routes, operating at very high signal levels to increase the signal to noise ratio, that were added together to form the detector. At the centre frequency of the discriminator the output

from the system was zero volts. The system was designed to work at impedance levels well above 50 ohms which ensured that the following LNA was easily designed, and all the circuits used large voltage swings to maximize the signal to noise ratio.

Since the signal has been converted to a low IF, a number of other possibilities are practical as well, including the use of a delay line discriminator. Again because the signal frequencies are lower, the signal levels and impedances can be increased substantially to improve the signal to noise ratio of the discriminator.

As with the Delay Line Discriminator it is important to remember that FM noise is being measured rather than phase noise, and the user needs to convert between FM noise and phase noise. The system measures both sides of the phase noise, so it is inherently displaying the average SSB phase noise. AM noise is stripped from the carrier and does not contribute to the measurement.

In principle, the system behaves in a similar way to the delay line discriminator method, but it does have some substantial advantages for some applications. In particular, since the discriminator operates at an IF, an amplitude limiter can be used to ensure the amplitude of the signal into the discriminator is always the same. Provided the input level is above a certain threshold (in the system the author used it could be used for RF levels as low as -30 dBm), the sensitivity of the discriminator can be independent of RF input level.

The method copes very well with drifting oscillators without the use of a PLL since the oscillator drift simply results in a change of DC output voltage from the discriminator. Provided the discriminator has been designed with a linear frequency versus output response over the used part of the DC output voltage, the response is independent of the precise frequency of the IF.

The system the author used included a filter to measure the residual FM of a signal source in a 300 Hz to 3 kHz bandwidth (a common signal generator specification parameter) and a tuned voltmeter to measure phase noise at 20 kHz offset to give a fast measurement of these two prime signal generator parameters¹². The development of VCOs, and even high frequency crystal oscillators, is considerably accelerated using this method since the operating conditions of the oscillator (supply voltages, bias currents, and oscillator tapping points) can be altered and the results displayed in “real” time with no concern about the calibration of the system changing due to changed RF levels.

The frequency down-conversion process to an IF can generate intermodulation products in the mixer that are at known frequencies. For higher RF frequencies, the products have a high order and change very rapidly if the frequency of one of the sources is changed, allowing them to be easily identified as artefacts of the conversion process. The high order of the intermodulation products can also spread their spectrum out significantly (making almost noise like humps), but as the input frequencies (and the intermodulation order) are raised, the intermodulation levels rapidly attenuate and become immeasurable. The effects are rarely measurable for frequencies above 100 MHz when using a 1.5 MHz IF.

¹² In principle this can be done for most of the systems, but this method is the only one where the output is always a measure of the wanted parameter, regardless of the signal level applied to the system

The performance of an FM discriminator system is limited by the noise figure of the amplifiers and limiters which recover the signal from the output of the mixer, and by the performance of the discriminator itself. In the case of the system previously described, the discriminator consisted largely of passive components, which exhibited very good noise characteristics, and the very high signal levels maximized the Hz/Volt at the output of the discriminator. Performance tends to be controlled by the slope of the discriminator, and it is for this reason that two band pass filters were used to allow a compromise between sensitivity and measurement bandwidth.

A high performance FM discriminator is capable of measuring very low levels of phase noise. The above system was capable of measuring residual phase noise of -170 dBc/Hz at 20 kHz offset and had a residual FM of <0.003 Hz in a 300 Hz to 3 kHz bandwidth.

6.1 Calibration of FM discriminator

Calibration of the system is very straightforward since the system sensitivity is independent of the input drive level to the frequency conversion mixer. Once a system has been constructed, the calibration factors are constants that can be allowed for by periodic (6 monthly) calibration checks. There is no requirement to perform a calibration before each test, a factor that can considerably increase the measurement speed

Calibration is typically performed by making one (or both) of the sources a signal generator with calibrated amounts of narrow band FM or phase modulation. Any switched amplifiers required to increase the signal levels will need to be calibrated. Provided that the system has been designed with stable components however, all the calibration factors are relatively time, frequency and RF level independent.

7 Digitizer Measurements

The previous methods described all assumed the use of bench instruments to do the spectral analysis of the output from the test system. However, the speed and performance of Analog to Digital Converters (ADC) has been advancing at rapid pace, driven largely by the demands of new radio receiver technologies (GSM, CDMA) and data acquisition systems.

The performance of ADCs is likely to change the balance of what can be achieved with bench instruments and what can be achieved by modular instruments. Modern spectrum analyzers already use high speed digitizers in their IF systems, but the design cycle for these instruments is long compared to the pace of converter technology development. At least one commercial phase noise measuring system is based on the use of modular instruments, but they are proprietary designs that will not necessarily keep up with the pace of converter technology development and the measurement method is focused on Delay Line Discriminator and the Quadrature methods.

An increasingly used route for phase noise measurements will be the use of high speed digitizers in open standard formats (such as PXI13, cPCI, PCI, VXI) which allow digital data to be rapidly exported to computing hardware that can perform analysis of the digitized data to extract phase noise information. Tools capable of performing (for example) FFT analysis of data files are available in a number of commercially available software tools.

The principles of the measurement will not be much different to those already described, but some variations may become attractive.

For low frequency applications, some sources may be tested by directly digitizing the source. It requires the clocks and circuits inside the digitizer to match or exceed the device under test. The converter used must have sufficient resolution to lower the quantization noise below the noise to be measured. Essentially the method used is the same as that previously described for a spectrum analyzer. However, the need for a swept frequency source has been removed. Functionally it is the equivalent of directly using a bench FFT analyzer but with more possibilities on the analog to digital converter hardware and the using third party software to provide the analysis of the results.




Figure 23: Phase noise measurement using direct digitization

A more attractive route is for higher frequency applications. The noise of the source under test is more likely to be higher than that of the digitizer, and converting the signal from a high frequency to an IF using a high quality LO and then digitizing the IF will improve the ability

¹³ For more information on the PXI standard a book, PXI mate, is available free from the Pickering Interfaces web site www.pickeringtest.com

to do higher performance measurements. Lower IF's allow the signal to be digitized with more resolution or using highly over-sampled clocking speeds to spread the quantization noise over a wider bandwidth.

Once the signal has been frequency down converted, it opens other possibilities for the analysis route to be used. The signal can be digitally demodulated for phase or frequency noise with none of the linearity issues likely to be introduced by the mixer/phase detector in other methods. The FM discriminator already described can be emulated with digital hardware rather than analog hardware, though for the foreseeable future, it is likely that analog hardware will exceed the capability of digital hardware.




Figure 24: Digitization with down conversion

The data files can be analyzed to separately extract phase and amplitude noise measurements.

Instead of measuring phase noise at fixed offsets, or as an offset versus noise display, noise effects can be windowed to reflect the impact the noise will have on the system (e.g. as a residual FM in an audio band, or impact on a digital modulation EVM).

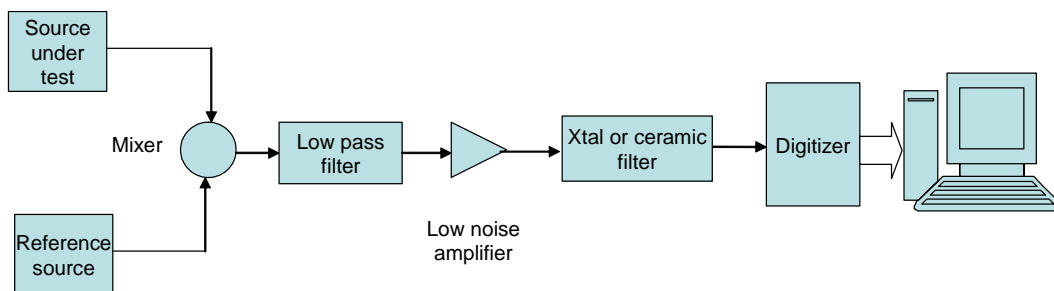


Figure 25: Digitization with down conversion and carrier removal

Both these methods rely on digitizing the signal with the carrier present (even if it is at an IF). For applications where close to carrier phase noise measurements are not required, the same method as described for spectrum analyzers can be used, with band stop filters in the IF to remove the carrier. As with the spectrum analyzer version, the performance required of the digitizer is considerably eased since the carrier has been removed during the critical measurement phase. It makes it more likely that a high performance source can be measured.

A useful future development of digitally based phase noise measurements may be the possibility of performing measurements in a synchronized way using combinations of three sources. If the three sources are mixed together in three separate mixers using the Quadrature

Method, then it may become possible to run synchronized measurements using three digitizers and uniquely extract the noise performance of each using correlation techniques.

7.1 Calibration of Direct Digital Measurement

The methods of calibration of the direct digital measurement methods are largely the same as for the analog methods they emulate.

The use of commercially available software tools is preferable to the use of “home brewed” software since the cost of proving the software can be very high compared to the cost of acquiring the software tools which may be of value to a wider range of applications.

8 Common problems with phase noise measurement

8.1 Reference oscillator

Finding a suitable LO or Reference Oscillator to test a source against can be problematic for high performance measurements. It is hard to measure the performance of a device when the noise from the test system is similar to the unit under test.

For measurements where the performance requirements are more modest, a low noise signal generator can be a valuable tool. Signal generators are available from a number of manufacturers and it can be hard to make a choice since they all have their strengths and weaknesses. A signal generator which has world-class performance at 10 kHz offset may be less effective close to the carrier, or at high offsets, for example. If a signal generator is chosen as the reference oscillator it must have the performance required at all the offsets which are likely to be required. In some cases, a user may require the use of more than one type of signal generator, making it an expensive investment in reference sources.

For some applications, it may be preferable to use a second sample of the device under test as the reference oscillator. The user can assume that both sources have the same performance and that the noise from them is not correlated. The performance of each of the sources can be approximated to be 3 dB better than the measured result

Note: This assumption can be suspect in the case of spurious signals. The two sources may have similar levels of spurious generated by a common mechanism. Depending on the phase of the signals, they can add or (more worryingly) cancel in the measuring system. This can lead to confusion in interpreting the true signal levels, especially when the user is trying to improve the performance of the source.

An alternative strategy to using two samples and assuming the performance is the same for each, is to use three samples of the product and measure the three in different combinations. Since three measurements are possible (A+B, B+C, A+C) and there are three unknowns (A,B,C), an estimate of the performance of each source can be made at particular offsets.

This 3 way measurement cannot be used reliably for measurements where an IF based measurement method is used, since the assumption is that the noise does not change between measurements. The source under test also has to be tunable to the offset frequency of the IF used. If it can tune to a suitable offset then a 6-way measurement is possible to extract an approximation for each source.

8.2 Injection locking

Injection Locking is a common problem when two sources of the same frequency are used in a system (in some systems injection locking is deliberately encouraged where the addition of a PLL would incur cost).

An oscillator is very susceptible to signals from external sources that are close to its operating frequency (how close is related to the Q of the oscillator). As the external source comes closer, its impact becomes more pronounced until the oscillator locks to the same frequency, in much the same way as a PLL behaves. When using the Quadrature Method of measuring

phase noise, interaction between the two sources used will usually encourage some sort of injection locking behavior that can be noticed unless the effect is swamped by other circuits. Injection locking effects can be reduced by ensuring that power supplies are well decoupled and isolated at RF, as well as at lower frequencies, and that the RF amplifiers in the system have high reverse isolation.

The effects of injection locking can be very variable and can cause some variation in close to carrier noise measurements. When using a PLL with the Quadrature Method to maintain phase quadrature of the signals, the BW of the PLL should be greater than the potential injection locking characteristics.

8.3 Reference oscillator locking

Besides the problems of injection locking, there are other risks from unintentional PLL's that are less obvious, concerning the use of frequency standards.

If a user is attempting to measure the phase noise characteristics of a synthesized frequency source, the frequency standard of the reference and the source under test is frequently made common so that getting phase quadrature stability is less problematic – for the Quadrature Method the external phase locked loop can be removed.

The connecting of the standards together can cause errors in the close to carrier phase noise (and sometimes out to offsets as large as a few hundred hertz). Any noise on the shared standard will appear on both the output signals and should cancel in the phase noise measuring system. The measurement to be made may require that this noise is included in the results, or the specification may require the measurement of just the “added” noise. There is an additional complication that the slave device will almost certainly have its own frequency standard conditioning circuits (perhaps another PLL or a crystal filter) to “clean up” the external standard. The slave synthesizer may even have a frequency standard that is better than the one being used at certain offsets, so the shared standard could make the phase noise appear to be worse at some offsets than it really is (it does not follow that the phase noise of a high stability standard is better at all offsets than a lesser stability standard).

It is essential that the consequences of using a shared standard are understood if there is a risk they could affect the measurement that is required to be made by the system. If the measurement is to include the frequency standard noise, it is best not to lock the source under test and the reference together and to check if there is any discernible tendency to injection lock.

8.4 Mixer linearity and errors

Throughout this guide, mixers have frequently been used as phase detectors to measure either phase noise or frequency noise. The mixers may not be ideal devices however. Mixers have also been used to provide frequency conversion to an IF.

The guide has generally assumed that the mixer is used with the recommended level of LO drive and a level into the RF port which allows the mixer to behave in a substantially linear way. The IF port is assumed to be DC coupled. When used in this way, if the LO inputs and the RF inputs are separated by an offset frequency, they will generate a low frequency sine

wave at their output. If the input frequencies are the same, the output is a function of a cosine of the phase difference between the two signals, the peak level being the same amplitude as the peak signal seen with a frequency difference between the two ports.

If a mixer is driven from a source impedance, or is loaded by a load impedance different from that which it was designed for (usually 50 ohms) it can exhibit unusual frequency response errors. The need to match the impedance extends to ensuring a reasonable match at the sum output frequencies from the IF port if reflections from this port are not to re-mix in the mixer. The IF port is usually the most sensitive to these errors; both the low frequency and high frequency signals it generates should be correctly terminated. The low pass filter that is usually used to remove the sum component should be designed with a diplexer to correctly terminate it.

The most critical matching problem occurs if the mixer is driven hard on both the LO and the RF port in order to improve the signal to noise ratio at the output of the mixer. In these circumstances, square waves are being generated in the mixer which are out of phase with the quadrature signals set up to the mixer. The output from the mixer in this case is a triangular wave with the phase difference, but the remixing of signals from reflections at the input ports can cause the triangular wave to have a non-linear slope that disrupts any attempt to calibrate the system.

If the input signal level to the RF port is low, it can confuse attempts to set phase quadrature at the mixer. With no signal applied to the RF port of a mixer, the output can exhibit a DC offset with measurable value (perhaps a few millivolts or more). The DC offset is inherent in the design of the mixer (usually reflecting slight imbalances in double balanced mixers) or may be partly due to input signal harmonics. To achieve phase quadrature the signal level at the mixer output may need to be set to this same DC offset. If possible, it is best to undertake measurements with (say) at least 50 mV of peak output voltage from the mixer to avoid errors¹⁴.

When mixers are used to frequency down convert a signal to an IF, they will introduce intermodulation products. These signals will appear as false spurious signals on the phase noise plots. Their level will be highly dependent on drive levels to the mixer, the type of mixer and the order of the intermodulation products that fall within the frequency band of interest. The order of the intermodulation product is given by the sum of the harmonic values of the signals that give rise to them (e.g. if the intermodulation product arises from the 5th harmonic of one signal and the 4th harmonic of another then it is a 9th order product). The higher the order of the product, the lower the level of the intermodulation signal generated. Intermodulation products involving the odd order harmonics of the LO are more problematic than those involving only even harmonics. The frequency of the intermodulation product also changes rapidly when one of the input frequencies is changed – making it relatively easy to identify when signals are artifacts of the systems rather than genuine spurious by simply doing two sets of measurements with the IF frequency changed by a small % of the frequency. Since IF based solutions are generally tolerant of small changes in frequency, this is technically not a great problem.

The table below gives some indication of the performance that might be expected from a standard Class 1 (+7 dBm), Class 2 (+17 dBm) and Class 3 (+27 dBm) mixers for relatively

¹⁴ More information on this is provided in the section on measurement uncertainty

low order intermodulation products. The exact performance of a mixer used for frequency down conversion will vary according to its style and manufacturer but the information does give an indicative guide.

| Harmonic | Class RF Level | 0 | | | 1 | | | 2 | | | 3 | | | 4 | | | 5 | | |
|----------|----------------------|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
| | | 1 | 2 | 3 | 1 | 2 | 3 | 1 | 2 | 3 | 1 | 2 | 3 | 1 | 2 | 3 | 1 | 2 | 3 |
| | | | | | | | | | | | | | | | | | | | |
| 0 | -10 dBm | | | | 26 | 27 | 18 | 35 | 31 | 10 | 39 | 36 | 23 | 50 | 47 | 14 | 41 | 36 | 19 |
| | 0 dBm | | | | 36 | 39 | 29 | 45 | 42 | 20 | 52 | 46 | 32 | 63 | 58 | 24 | 45 | 37 | 29 |
| 1 | -10 dBm | 24 | 23 | 24 | 0 | 0 | 0 | 35 | 39 | 34 | 13 | 11 | 11 | 40 | 46 | 42 | 24 | 14 | 18 |
| | 0 dBm | 25 | 25 | 24 | 0 | 0 | 0 | 39 | 39 | 35 | 13 | 11 | 11 | 45 | 50 | 42 | 22 | 16 | 19 |
| 2 | -10 dBm | 73 | 86 | 73 | 73 | 75 | 83 | 74 | 84 | 75 | 70 | 75 | 79 | 71 | 86 | 80 | 64 | 74 | 80 |
| | 0 dBm | 69 | 68 | 64 | 72 | 67 | 71 | 79 | 76 | 62 | 67 | 67 | 70 | 75 | 80 | 63 | 66 | 66 | 70 |
| 3 | -10 dBm | 67 | 87 | >90 | 64 | 77 | >90 | 69 | 87 | >90 | 50 | 78 | >90 | 77 | >90 | >90 | 47 | 75 | >90 |
| | 0 dBm | 51 | 63 | 81 | 49 | 58 | 73 | 53 | 65 | 85 | 51 | 60 | 69 | 55 | 65 | 85 | 48 | 55 | 68 |
| 4 | -10 dBm | 86 | >90 | >90 | >90 | >90 | >90 | 86 | >90 | >90 | 88 | >90 | >90 | 88 | >90 | >90 | 85 | >90 | >90 |
| | 0 dBm | 80 | 96 | 83 | 79 | 80 | >90 | 82 | 96 | >90 | 77 | 80 | >90 | 82 | >90 | >90 | 76 | 82 | >90 |
| 5 | -10 dBm | >90 | >90 | >90 | 80 | >90 | >90 | >90 | >90 | >90 | 71 | >90 | >90 | >90 | >90 | >90 | 68 | >90 | >90 |
| | 0 dBm | 72 | >90 | >90 | 70 | 73 | >90 | >90 | 87 | >90 | 52 | 72 | >90 | 77 | 88 | >90 | 46 | 66 | >90 |
| 6 | -10 dBm | >90 | >90 | >90 | >90 | >90 | >90 | >90 | >90 | >90 | >90 | >90 | >90 | >90 | >90 | >90 | >90 | >90 | >90 |
| | 0 dBm | >90 | >90 | >90 | 86 | >90 | >90 | >90 | >90 | >90 | >90 | >90 | >90 | >90 | >90 | >90 | 84 | >90 | >90 |

As can be seen from the table spurious levels become negligible for intermodulation products beyond the order of 10 for most purposes, allowing 10.7 MHz IF systems to be used satisfactorily for frequencies above approximately 70 MHz.

8.5 Power Supplies

All sources are likely to be capable of being affected by their power supply noise. It is essential that when a source is measured, it is being powered by a supply that is representative of what will be used in its final application. The noise on the power supply will be converted from a time varying voltage to a phase modulation signal by non-linearity in the oscillator. In many oscillator designs the optimum operating point, particularly for close to carrier phase noise, is often close to the point at which the oscillator frequency is least likely to change with operating voltage or current. Large dependence of frequency on oscillator voltage can be an indication that power supply noise is a contributing factor to the noise performance of the oscillator.

Where oscillators are required to have good far out noise, the oscillator may be operated at a level where the close to carrier noise has been deliberately allowed to degrade in favour of the critical noise parameter, and that can make the source particularly vulnerable to power supply noise.

Many modern power supplies use switched mode designs that contain significant spurious signals that need to be removed if the true performance of the source is to be measured.

8.6 Detectors and filters

Phase noise is a measurement of relative power of the carrier and the noise at a given offset. If the signals are measured with anything other than a power detector there is the potential for calibration errors to be introduced.

Usually these errors will not be very large, but when a designer is struggling to achieve the last few dB of phase noise improvement in the source under test it can be a crucial issue.

The final detector usually resides in a spectrum analyzer. If the spectrum analyzer is an entirely digital device (a digitizer or an FFT Analyzer), then it is simply a question of ensuring that the software processes the digital information in the required way (which is usually the case).

Older style spectrum analyzers use analog detectors on their final IF to measure the signal. These are never power detectors and they measure noise signals and coherent signals (sine wave) with different levels for the same level of power. They also process the signal through logarithmic amplifiers to compress the dynamic range of the signal that is being measured. Again this influences the detected level measured for coherent and noise signals.

For an analog spectrum analyzer two correction factors are usually quoted¹⁵:

Detector correction for noise. When the analyzer is calibrated, it is assumed that the signals to be measured are coherent (not noise). The detector is typically peak responding, so the level is assumed to be 3 dB lower (the peak of a sine wave compared to its RMS value). Noise, however, is detected differently, and appears to be 1.05 dB in error.

Detector Correction for Log scaling. Users almost always use a spectrum analyzer on a logarithmic vertical scale. The averaging process after the logarithmic detector causes a 2.5 dB level error.

The net result of these two effects is that noise on an analog spectrum analyzer typically reads 1.45 dB below its true level.

More modern spectrum analyzers can have digitizers that capture the signal and can be used to measure the signal at the final IF. These can correctly display the levels for both noise and spurious signals if they are set to use a “true RMS” or “power” detector. The user has to wary though; the spectrum analyzer can be set so that it seems to behave in the same way for both the analog and digital detection methods.

Many modern spectrum analyzers may use both analog and digital detection mechanisms. They use digital filters only for narrow filter bandwidths and analog filters for the wider bandwidths. In this case, the user will have to consult the operating manual or the manufacturer’s application notes to discover how the instrument handles noise signals.

¹⁵ For an explanation of these factors a reader can consult Application Note 150 “Spectrum Analysis Basics” on the Agilent web site at www.agilent.com

In general, spectrum analyzers are calibrated to display the relative levels of coherent (sine wave) signals, rather than noise signals, and that is how they are calibrated by calibration centers in service.

8.7 Stitching errors

It is quite common to see a test system produce a phase noise plot where the phase noise suddenly changes in level by a few dB's. It is usually a sign that the calibration of the plot is incorrect and the results cannot be relied on.

The problem usually arises because the equipment in the system changes bandwidth or some other critical setting (such as an amplifier or PLL setting) in the system that has an impact on the correction values applied to a measurement. Sometimes it can indicate that part of the test system was overloaded during one set of measurements.

As an example, in the Quadrature Method, the PLL bandwidth may have been changed, the spectrum analyzer settings altered (input attenuation, filter bandwidth, IF gain); and the correction information applied has then resulted in the phase noise at the offset frequency, where the stitching error occurs, being measured differently. Though it is possible that the noise may have changed while the system settings were reset, it is likely that one or both the measurements are in error, and the measurement must be redone after identifying the likely cause.

These sorts of errors are more likely to be found in the measurements based on phase demodulation than frequency demodulation, because there are fewer range and calibration changes involved.

8.8 Spurious noise identification and display

In all phase noise measurements, the noise is measured in a bandwidth which is typically not 1 Hz. Part of the correction data that is applied is to convert the noise in the measured bandwidth (e.g. 10 Hz) and display the noise in a 1 Hz bandwidth (in this case a correction of 10 dB). However, if the signal being measured is not noise but instead contains a coherent (spurious) signal the level of this signal (assuming it is a narrowly defined frequency) is the same irrespective of what bandwidth it is measured in.

If the signal at a given offset is dominated by a spurious signal, the noise bandwidth correction cannot be applied. Instead, the signal should be shown as a line with its amplitude equal to the measured value in dBc. That is also the level at which a spectrum analyzer will measure the signal.

Some software tools in the equipment may attempt to identify noise and spurious signals and correct the reading automatically. Unfortunately, they do not always succeed, especially when the signal is only marginally above the surrounding noise levels. The algorithms may be helped by narrowing the measurement bandwidth (at the expense of measurement time), because this lowers the level of the measured noise compared to the spurious signal, making it easier for the software to distinguish between noise and a spurious signal for a given spurious signal level (note that this only moves the problem to a lower level, it never eliminates the problem).

Phase noise plots with spurious signals in them often lead to user concern at what are perceived to be high levels of spurious in the source; concerns that can be unfounded. Taking a specific example a source may have a spurious at 100 kHz offset that is -100 dBc, and the phase noise may be -140 dBc/Hz. If a phase noise plot shows the signal spectrum with the spurious correctly displayed at -100 dBc and the noise at -140 dBc/Hz, it appears to be some 40 dB above the noise floor, despite the fact that such a signal is almost immeasurable on a spectrum analyzer. If the same plot is shown with noise bandwidth correction applied to all the information (including being erroneously applied to the spurious) and the signal measured in a 1 kHz bandwidth then the spurious will only appear to be 10 dB above the noise. Be wary of what you are looking at on the plot!

Increasingly in a digital world, there is a tendency for more spurious signals to be present from high-speed clocks, their subharmonics and their harmonics. This raises the possibility that some “noise” signals will not be noise, but a collection of two or even a great many closely separated spurious signals. It is not uncommon for these to be mistaken as noise, especially if their level has been shaped by power supply filtering or other effects.

8.9 Dithered clocks

There has been a tendency in recent years to use dithered clocks in systems. These clocks have deliberately introduced phase modulation so that their spectrum is spread out. In some cases, such as for clocking ADC's, this has real systems performance benefits. In other cases, it may be simply to reduce the emission levels measured by the defined conditions in EMC conformance tests. Distinguishing these signals can be problematic since they appear to be noise-like, but in reality are coherent signals. The amplitude distribution of these signals will still be more like a sine wave rather Gaussian noise, so it can invalidate detector correction factors. For phase noise measurements it may be best to disable the dither signal while measurements are being made that are intended to help find true noise levels.

8.10 Software correction artifacts

Besides the issue of identifying spurious signals and noise, some measuring instruments use more subtle techniques to expand the displayed dynamic range or to make noise signals appear to be more realistic. These artifacts can be a problem for phase noise measurements, leading to errors in interpreting the actual signal levels measured. As an example, a software correction method seen on at least one spectrum analyzer results in ambiguous noise measurements when the instrument is used close to its “known” residual noise floor – attempting to correct the displayed results for a known limitation in the product. The linearity of the spectrum analyzer is best assured by making sure that it is operated near the centre of the its “display” range and as far away from its residual noise floor as possible.

This type of problem should be less prevalent on digitizer products where the user has access to “uncorrected” data and has better control over the software algorithms used.

8.11 Sideband selectivity

Phase noise measuring systems often have problems from the mixing of different sidebands together. The most common is the problem of the addition of the upper sideband and lower

sideband noise. This applies to all but the direct spectrum measurement systems (spectrum analyzer or digitizer measurements). Frequency and phase discriminators respond to both the upper and the lower sideband noise and consequently display the sum of these two noise components. They assume that the spectrum is symmetric, which is not always the case. The displayed noise is the average of the two sidebands.

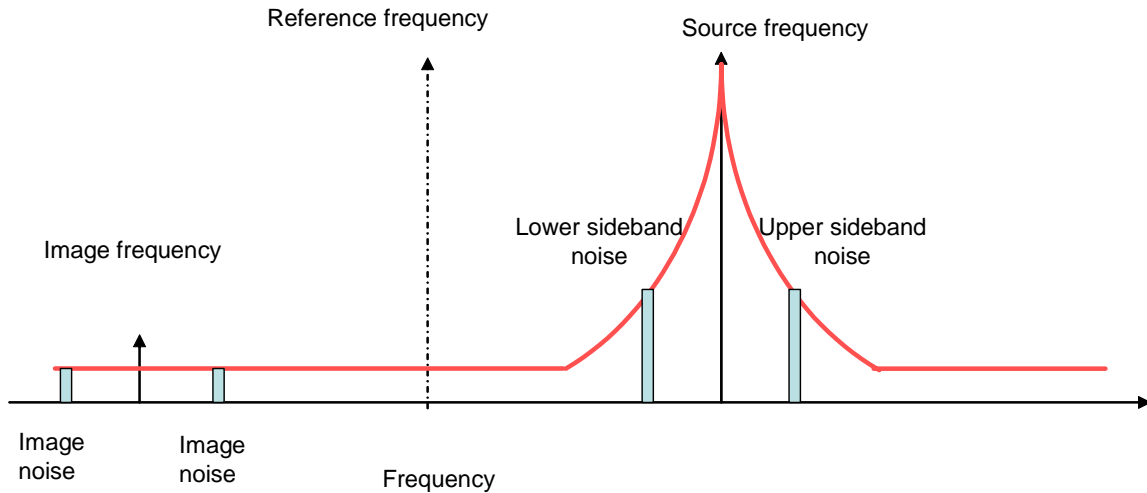


Figure 26: Showing the addition of unwanted noise in an IF based system

On IF based systems there is a further complication. Signals (including noise) at the image frequency are mixed with the wanted noise signals and add to the measured results. Usually this is not a problem since phase noise falls with increasing offset frequency so the added noise from this process can be neglected. However, a few systems may have a rising noise profile at this offset (e.g. induced by a wide bandwidth PLL) and may give unexpected answers. In systems where the phase noise profile is flat from the measured frequency to the image frequency (e.g. where it is limited by floor noise amplifiers) the noise level measured will be 3 dB worse than the actual level

8.12 AM Noise rejection

By definition, phase noise measurements should not include the contribution of amplitude noise in the measurement. However, not all the measurement methods exclude the amplitude component from the measurement.

Measurements based on the use of a spectrum analyzer, whether at IF or directly on the RF signal, will include the amplitude component. If the amplitude component of the noise is high, it may give a misleading answer since the measurement then adds both noise sources together.

Measurement methods that use a mixer as a phase detector by having two signals in quadrature at its input tend to reject the amplitude noise component since the mixer is just detecting phase differences. Any error in the phase at the input to the mixer will cause some of the AM component to appear in the measured output.

The FM discriminator method strips the amplitude noise from the carrier when the IF signal goes through the limiter provided the limiter has been designed so that it does not convert amplitude signals into phase variations.

Some methods can be set to measure the AM component separately from the phase component. The Quadrature Method can be adjusted to provide in phase signals to the mixer so it behaves as an amplitude detector. The output from the mixer is tuned for maximum magnitude of voltage and the DC pedestal has to be removed by AC coupling the output to the following detectors. Digitizer approaches can be used to simultaneously extract the amplitude and phase component.

In general amplitude noise is likely to be a factor only where the signals being measured have been processed through amplitude control systems that have relatively poor signal to noise ratios, and the problems are usually confined to the higher offset frequencies rather than close to carrier signals.

9 Comparison of the Methods

Drawing up comparisons of methods of measurement is always a sensitive subject. Commercially available phase noise measurement systems will go to great lengths to improve the performance of a particular method, and consequently may be able to take measures that partially overcome (or disguise) the limitations.

However, commercially available systems can be relatively expensive to purchase, and many users with a need to perform phase noise measurements cannot justify the purchase of an item that may only get limited use. Instead, systems are put together from equipment already available in the laboratory. These systems may also be far better focused on a particular measurement requirement, and in particular may be able to take advantage of sources or other equipment not readily supported in a commercial system. A purpose built system can also be much faster at performing specific tests required in a manufacturing or quality control environment.

The following tables are intended to provide objective guides to a user who has to construct his own system. It is based on the author's experience of both proprietary and commercial phase noise systems. There is no guarantee that an inexperienced user will achieve these results, or that an experienced engineer will agree with the summary.

By application

| | Spectrum Analyzer | | Delay Line Discriminator | Quadrature | | FM Discriminator | Digitizer | |
|--|--|--|--|---------------------------------|--------------------------------------|--------------------------------------|------------------------------------|--|
| | Direct | IF and filter | | No PLL | With PLL | | Direct | IF and filter |
| Drifting oscillators | Only at large offsets and very limited performance | Marginal, good if drift is reasonably low | Very Good | Impossible | Good with care, but more complicated | Very Good | Very good with additional software | Marginal |
| Oscillators with large low frequency noise | Good but very limited performance | Very Good, but will not measure close to carrier | Very Good | Impossible | Good with care, but more complicated | Very Good | Very Good | Very Good |
| Microwave sources | Can be good within the performance limits | Can be good within the performance limits | Gets harder to handle the frequencies | Probably impossible | Very Good | Good within the limited offset range | Impossible | Very Good |
| Crystal Oscillators | Impossible | Performance unlikely to be good enough | Performance unlikely to be good enough | Very Good | Very Good | Performance may not be good enough | Performance not good enough | Very good, but limited on bandwidth |
| Large offset noise | Very limited performance | Very Good | Limited suitability | Very Good | Very Good | Not suited | Very limited performance | Can be good, but limited suitability |
| Real time iterative improvement of oscillators | Good but limited to low performance | Limited by shifting calibration | Limited by shifting calibration | Limited by shifting calibration | Limited by shifting calibration | Very Good | Very Good | Very Good |
| Close to carrier noise measurements <20 Hz | Not suited | Not suited | Not suited | Very Good | Very Good | Not suited | Very Good | Can be very good |
| 1 KHz phase noise | Very limited | Limited | Good | Very Good | Very Good | Very Good | Very Good | Can be good with use of narrow filters |
| Asymmetric spectrum | Good with limited performance | Good with limited performance | Impossible | Impossible | Impossible | Impossible | Potentially very good | Potentially very good |

Expected Performance Limitations

| Offset Frequency | Spectrum Analyzer | | Delay Line Discriminator | Quadrature | | FM Discriminator | Digitizer | |
|------------------|--|--|---|--|--|---|---|---------------------|
| | Direct | IF and filter | | No PLL | With PLL | | Direct | IF and filter |
| 0.1 Hz | Unsuited | Unsuited | Unsuited | Limited by reference source | Limited by PLL calibration issues | Unsuited | Limited by quantization and clock noise | Unsuited |
| 1 Hz | Unsuited | Unsuited | Unsuited | Limited by reference source | Limited by PLL calibration issues | Unsuited | Limited by quantization and clock noise | Unsuited |
| 10 Hz | Unsuited | Unsuited | Unsuited | <-130 dBc/Hz Limited by reference source | <-130 dBc/Hz | Unsuited (typically usable at 20 Hz) | Limited by quantization and clock noise | Unsuited |
| 100 Hz | Unsuited | Unsuited | -60 dBc/Hz | Limited by reference source | Limited by reference source | <-110 dBc/Hz | Limited by quantization and clock noise | Unsuited |
| 1 kHz | -90 dBc/Hz Limited by LO noise | Unsuited | -95 dBc/Hz | Limited by reference source | Limited by reference source | <-140 dBc/Hz | Limited by quantization and clock noise | <-160 dBc/Hz |
| 10 kHz | -115 dBc/Hz Limited by LO noise | <-155dBc/Hz Limited by filter selectivity | -125 dBc/Hz | <-160 dBc/Hz Limited by thermal and reference performance | <-160 dBc/Hz Limited by thermal and reference performance | <-160 dBc/Hz | Limited by quantization and clock noise | <-160 dBc/Hz |
| 100 kHz | -120 dBc/Hz Limited by LO noise | <-160 dBc/Hz | -145 dBc/Hz | <-160 dBc/Hz Limited by thermal and reference performance | <-160 dBc/Hz Limited by thermal and reference performance | <-160 dBc/Hz Limited by thermal noise and noise figure | Limited by quantization and clock noise | <-160 dBc/Hz |
| 1 MHz | -140 dBc/Hz Limited by LO noise and dynamic range | <-160 dBc/Hz | <-160 dBc/Hz | <-160 dBc/Hz Limited by thermal and reference performance | <-160 dBc/Hz Limited by thermal and reference performance | Unsuited | Limited by quantization noise | <-160 dBc/Hz |
| 10 MHz | -140 dBc/Hz Limited by LO noise and dynamic range | <-160 dBc/Hz | Unsuited (assumes 100 ns delay line is used) | <-160 dBc/Hz Limited by thermal and reference performance | <-160 dBc/Hz Limited by thermal and reference performance | Unsuited | Limited by quantization noise and BW | Probably not suited |

Overall Performance

| Spectrum Analyzer | | Delay Line Discriminator | Quadrature | | FM Discriminator | Digitizer | |
|-------------------|---------------|--------------------------|------------------------------------|------------------------------------|--------------------------------------|-----------------------|-----------------------|
| Direct | IF and filter | | No PLL | With PLL | | Direct | IF and filter |
| Low | Modest | Modest | Very Good Highest dynamic range | Very Good Highest dynamic range | Good But limited range of offsets | Potentially very good | Potentially very good |

Ease of Calibration

| Spectrum Analyzer | | Delay Line Discriminator | Quadrature | | FM Discriminator | Digitizer | |
|-------------------|--------------------|--------------------------|--------------------|---------------------|------------------|-----------|--------------------|
| Direct | IF and filter | | No PLL | With PLL | | Direct | IF and filter |
| Good | Some complications | Some complications | Some complications | Can be very complex | Very Good | Very Good | Some complications |

What it measures

| Spectrum Analyzer | | Delay Line Discriminator | Quadrature | | FM Discriminator | Digitizer | |
|--------------------------|-------------------------------------|-----------------------------------|--------------------------------------|--------------------------------------|---|--------------------------|-------------------------------------|
| Direct | IF and filter | | No PLL | With PLL | | Direct | IF and filter |
| AM+Phase Single sided | AM+Phase plus image Single sided | FM Averaged of upper and lower | Phase Averaged of upper and lower | Phase Averaged of upper and lower | FM Plus image Averaged of upper and lower | FM or AM single sided | FM or AM single sided plus image |

Speed

| Spectrum Analyzer | | Delay Line Discriminator | Quadrature | | FM Discriminator | Digitizer | |
|------------------------------|---|--------------------------|--------------|-----------------------|--------------------|--------------------|-----------------|
| Direct | IF and filter | | No PLL | With PLL | | Direct | IF and filter |
| Fast for what it can measure | Reasonably fast for what it can measure | Modest speed | Modest speed | Can be time consuming | Fast and real time | Fast and real time | Reasonably fast |

Reference Oscillator Requirements

| Spectrum Analyzer | | Delay Line Discriminator | Quadrature | | FM Discriminator | Digitizer | |
|----------------------------|------------------|--------------------------|--|--|------------------|----------------------------------|------------------|
| Direct | IF and filter | | No PLL | With PLL | | Direct | IF and filter |
| None, embedded in analyzer | Offset reference | None | At same frequency as source under test | At same frequency as source under test | Offset reference | None, depends on converter clock | Offset reference |

Other issues or complications

| Spectrum Analyzer | | Delay Line Discriminator | Quadrature | | FM Discriminator | Digitizer | |
|---------------------|--|---|---|--|---|---|--|
| Direct | IF and filter | | No PLL | With PLL | | Direct | IF and filter |
| Limited performance | No commercially available solution Image noise adds | Hard to automate in user constructed systems Can measure with some deliberate FM present | Injection locking Low frequency noise overload | PLL calibration Low frequency noise overload Need voltage tuning | No commercially available solution Image noise adds Can measure with some deliberate FM present | No commercially available solution Can measure with some deliberate FM present | No commercially available solution Image noise adds |

10 Estimating Uncertainty

The measurement methods for measuring phase noise can be subject to large error bands. Commercial solutions typically offer a specification of 2 dB on the noise measurement, but usually there is no breakdown on how this arrived at, or how it may be affected by external factors.

Since there is considerable diversity in how, and under what conditions the measurement is done, there is no unique answer as to how accurate a measurement is. It is probable that, in the long-term, digitally based solutions will offer the best accuracy.

The following is intended to give some guidance on estimating the uncertainty of the measuring instruments typically seen in phase noise measurement systems and to highlight the major sources of error. It also seeks to show how the uncertainty can be reduced.

One of the fundamental problems of phase noise measurement is that it is a ratiometric measurement that requires very high dynamic range to perform. The most effective methods therefore rely on removing the carrier (by filtering or phase/frequency detection), so removing the reference value for any measurements.

These errors make it difficult to assign a traceability figure to phase noise measurements. Often measurements of the same device will lead to differing answers, even when measured on the same system. It is not uncommon for instance to see “stitching” errors in the results where the test system changes settings to measure noise at different offset frequencies.

10.1 Spectrum Analyzer Uncertainty

Spectrum analyzers form an essential part of most phase noise measuring systems. The definition of a spectrum analyzer in this guide included both analogue and digital spectrum analyzers and the use of FFT analyzers.

In general FFT analyzers will produce the most accurate measurements since, like the “Digitizer” based measurements already described, the devices have relatively simple analogue front ends, a powerful ADC and an imbedded computing engine that derives a plot of the frequency dependency of the noise signal being measured. The digital nature of the instrument also allows measurements to be readily made with narrower (digital) filters and close-to-carrier results can be more easily gathered. However, not every laboratory will have an FFT analyzer, so conventional (digital or analogue) spectrum analyzers are commonly used. A test system may, for example, already include a conventional spectrum analyzer that is performing other measurement functions (spurious, harmonics) that an FFT analyzer cannot measure.

A conventional spectrum analyzer is a complex machine with a considerable number of analogue components, as was shown in the section on spectrum analyzer based measurements.

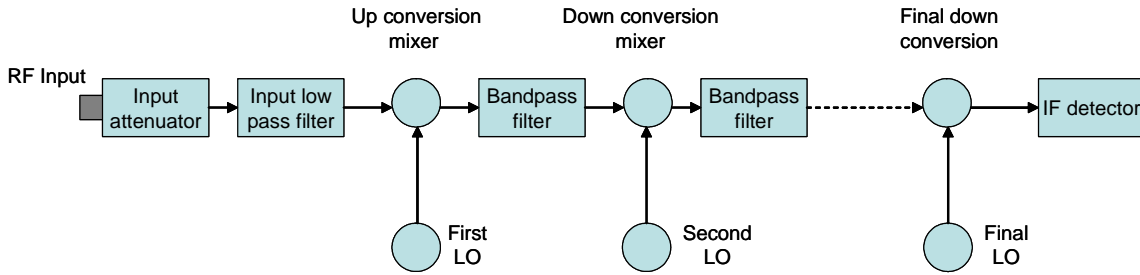


Figure 27: Spectrum analyzer block diagram

These analogue components can introduce considerable errors in the accuracy of the measurements they perform, particularly on the most commonly found relatively low cost instruments. The way the specification for a spectrum analyzer is written makes it hard to assign a single number to the errors introduced; the errors in the measured level come from a number of factors that the user has to take into account.

The following highlights the principle issues that may affect the overall uncertainty due to changes in the spectrum analyzer performance. It is assumed that some sort of calibration has been performed using the spectrum analyzer as a measuring device so that most of the errors that need to be identified are ratiometric. In general, absolute accuracy is not the most fundamental requirement, but in some cases (e.g. the FM Delay Line Discriminator Method, where the calibration method may produce an absolute measure of system sensitivity) it may be important.

The specification examples are taken from real spectrum analyzer data sheets, generally not the most expensive so not the best that are available, but certainly representative of the sort of spectrum analyzer that might be readily available to a typical user in a modestly equipped laboratory.

Input Attenuator. When the calibration is performed, the signal level to the spectrum analyzer is likely to be relatively high. The calibration may have to be performed with relatively high input attenuator settings that then need to be changed when the real measurements are done. Spectrum analyzers typically do not have a good noise figure, so using it at an optimum range for both the calibration cycle and the measurement cycle can be an important factor. The spectrum analyzer will have a specification for its input attenuator which might be of the form:

$$\pm 0.6 \text{ dB per } 10 \text{ dB step}$$

Changing the input attenuation by 20 dB could add up to 1.2 dB of uncertainty, though for most cases the error is likely to be rather less. FFT analyzers will typically have lower errors quoted because of their lower frequency range.

Input VSWR. The errors can be compounded by VSWR effects. The spectrum analyzer input attenuator performance is measured in a good 50 ohm system. If the system has a poor match then additional VSWR errors can arise. The specification can read something like:

$$< 1.5 \text{ for an input attenuation of } 10 \text{ dB or greater for frequencies } 30 \text{ Hz to } 2.9 \text{ GHz}$$

Each time the input attenuation is changed, the errors introduced by the attenuator can change because of the changing input VSWR. In reality, the largest errors are likely to be at high frequencies, but it is advisable to ensure the instrument is working in a good 50 ohm environment. Tables of the uncertainty caused by mismatches at each end of a 50 ohm cable are commonly published¹⁶.

IF Gain errors. Once the signal is in the spectrum analyzer IF system it will pass through a variety of switched gain amplifiers. If the gain of the IF is changed it will have an impact on the accuracy of the measurement. A typical spectrum analyzer specification might be something like:

IF Gain Uncertainty ± 1 dB for 0 to -80 dBm reference levels.

Consequently changing the IF gain (in other words reference level changes without a change in input level) from the calibration cycle to the measurement cycle can introduce this uncertainty.

Filter BW changes. Every time the filter BW is changed from the BW used in the measurement cycle a new filter with a different set of responses is introduced into the system. A spectrum analyzer specification might read something like:

Resolution BW switching uncertainty: ± 0.5 dB.

If the calibration cycle is performed on a 300 Hz filter BW setting, changing to narrower filters for close to carrier plots, or wider filters for faster high offset measurements has the potential to introduce an additional uncertainty.

Noise BW errors. For analogue based spectrum analyzers real filters will be imprecise on their stated equivalent noise BW. The filter responds mainly to signals within its pass band but has residual responses out of its bandwidth. The noise BW is not the same as the 3 dB bandwidth quoted in the data sheets. Where the filter is implemented with analogue components the shape and conversion factor may be different for different filters. The manual for the spectrum analyzer must be consulted to identify these changes.

A typical Agilent spectrum analyzer using a 10 kHz filter, for example, is quoted as having a noise BW that is between 1.05 and 1.13 the times the 3 dB BW of the filter. The ratio of the noise BW to 3 dB bandwidth of the filter creates a spread 0.32 dB in the BW correction figure (± 0.16 dB about the median value). In addition, the 3 dB BW may be in error from its nominal value, creating an additional systematic error.

Digital filters are likely to be better than analogue filters since their shape is largely described by software algorithms and (provided the window shape is not changed) they should at least only introduce a systematic error. For all these measurements, it is assumed that the noise density of the signal does not change significantly across the filter BW, so the BW of the measurement needs to be at least an order of magnitude narrower than the offset frequency at which it is measured.

¹⁶ Examples can be found in the Aeroflex or IFR RF Data Mate available through www.aeroflex.com

If the filters are measuring a noise signal whose frequency response is significant within the response of the filter (including the responses on the skirts of the filter) it may lead to additional uncertainty, especially if the shape of the filter is asymmetric. The only solution to this is to use a filter setting which is narrow compared to the rate of change of noise level with frequency.

Scale fidelity. The detector used in the spectrum analyzer has its own errors in detecting the actual signal level. The level errors at the top of the screen are different to those in the middle or near the bottom. A typical specification (in the logarithmic mode) might be of the form:

For residual BW of ≥ 300 Hz 0.1 dB/dB with a maximum error of ± 0.85 dB

The errors may be different according to BW. As the signal is moved from the level used on the screen in the calibration cycle, additional error is introduced which increases as the deviation from that displayed setting is increased. There is usually a “worst case value” given as well which is less than that which would occur from the specification. Again, digitally based systems are usually less prone to error than analogue based systems.

Detector Response. Unless the spectrum analyzer is using a true RMS responding measuring system, it will measure noise differently to coherent signals. The calibration process is usually based on the measurement of coherent signals (e.g. by offsetting the two frequencies to obtain a tone in the Quadrature Method). The correction is usually made by nominal correction factors provided in the user manual. Typically, an analogue spectrum analyzer using averaged logarithmic detectors will measure noise 1.45 dB lower than it actually is.¹⁷

Input frequency response. If a signal is applied to the front of spectrum analyzer, it will have frequency response errors. For most phase noise measurements, this response error has a small impact on the measurement if the instrument has a considerably wider BW than the signal being measured. However, digital systems (such as FFT analyzers) are more likely to have frequency response errors if their BW is not very large. Spectrum analyzers should be used with care at their lowest frequencies to avoid the AC coupling BW restrictions present on many models.

With the potential errors that have been described above a reader could be forgiven for wondering how any accurate measurements can be made of phase noise with a spectrum analyzer. However, modern spectrum analyzers with digital subsystems are considerably more accurate (and stable) than their older counterparts, and they often perform rather better in accuracy than the above would suggest¹⁸. For specific test systems the accuracy of the analyzer under the measurement conditions actually used can be tested by a calibration process to provide a revised (and better) uncertainty.

Digital instruments are in general much better than analogue instruments since some of the sources of error become virtually negligible.

¹⁷For an explanation of these factors a reader can consult Application Note 150 Spectrum Analysis Basics on the Agilent web site at www.agilent.com

¹⁸ See the Product Note Agilent PSA Performance Spectrum Analyzer Series Amplitude Accuracy on the Agilent web site www.agilent.com for an example of what can be achieved on a modern high performance analyzer

10.2 System Level Sensitivity

Most of the methods (the exceptions being the FM Discriminator and potentially Direct Digitizers) rely on a calibration cycle being performed before the measurement cycle to obtain a system sensitivity measurement. This cycle introduces its own uncertainties.

An obvious constraint is that the signal levels must be stable over the entire measurement period if accurate results are to be obtained. If the level changes by 0.2 dB over the measurement cycle, the uncertainty adds to the accuracy of the measurement.

There are additional sources of uncertainty introduced where mixers are used as phase detectors (whether in an FM Delay Line Discriminator or the Quadrature Method). In most cases, the calibration involves establishing the peak output voltage of the mixer and using this as a measure of the sensitivity of the system.

If the output of the system is assumed to have a sine wave dependency on the phase difference of the two signals there are some potential problems:

The output may have distortion in it (easily checked in the case of the Quadrature Method) that leads to the peak value not being an accurate reflection of the slope of the system at phase quadrature. Compression will lead to underestimation of the slope and give pessimistic results.

The detector may have some residual DC component when the linear input signal is not present. The maximum sensitivity may therefore not be with zero voltages at the output of the mixer. Since the peak voltage measurement indicates the maximum sensitivity, it is possible the system will perform a measurement with the mixer set to a lower sensitivity operating condition than is calculated. It can then give optimistic measurements. This can be compensated for, by measuring both the peak positive voltage and the peak negative voltage to derive the DC offset value. The system should then be operated at the DC offset voltage to minimize errors. The DC offset can be a secondary effect from distortion in the two input signals to the mixer, a high residual DC offset (compared to that specified for the mixer) might indicate additional input filtering is required for accurate measurement.

In other cases, the mixer may be driven hard to produce a triangular dependency on the phase difference of the two input signals. This opens other possibilities of error:

Poor VSWR at the inputs to the mixer may generate reflections that result in the mixer slope not being linear

The peak voltage may not be a good measure of the slope because of “rounding” of the waveform at the top and bottom.

Neither of these effects are easy to quantify, the user should make his own assessment of whether they are significant factors.

In some cases system uncertainty may be reducible by considering which signal drives which port of the mixer. If a mixer is driven in a linear way (to get a sinusoidal response) the LO port is driven harder than the Reference port. Variations in the level at the LO port (if the

signal is large enough) only have a secondary impact on the output sensitivity of the mixer whereas variations in the Reference port level have a direct impact.

The measurement conditions used in this guide have all assumed that a mixer used as a phase detector is adjusted so there is no DC voltage on the output. If the measurement is taken with a residual voltage present then there will be an error in the measurement (it will generally give an optimistic result).

The table below shows how sensitive the system is to an error assuming that the mixer produces a maximum output voltage of 1 Volt when the signals are in phase.

| Phase difference | 80 | 81 | 82 | 83 | 84 | 85 | 86 | 87 | 88 | 89 | 90 |
|------------------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|
| DC Output | 0.174 | 0.156 | 0.139 | 0.122 | 0.105 | 0.087 | 0.070 | 0.052 | 0.035 | 0.017 | 0.000 |
| Slope | 0.985 | 0.988 | 0.990 | 0.993 | 0.995 | 0.996 | 0.998 | 0.999 | 0.999 | 1.000 | 1.000 |
| Error dB | 0.13 | 0.11 | 0.08 | 0.06 | 0.05 | 0.03 | 0.02 | 0.01 | 0.01 | 0.00 | 0.00 |

It can be seen the system is not very sensitive to residual DC voltages being present, but for best results, it should be restricted to less than 10% of the peak output voltage.

The same table also shows that the peak phase excursions of the signal at the output of the mixer should be restricted to less than 0.1 radians (approximately 6 and 12% of the peak output voltage) if low frequency noise is not to cause the mixer to behave in a non-linear way and cause the generation of optimistic results.

Note that this table does not include any allowance for the residual DC outputs of the mixer mentioned above.

10.3 PLL Estimation

The Quadrature Method is particular often relies on a PLL to maintain phase quadrature at the mixer to provide a stable phase detector. If measurements of noise are to be made inside the PLL BW, it is essential to estimate the uncertainty of any correction factors applied. This will have to include some sort of estimate of the risks of injection locking (whose effects may be hard to predict) and the level variations after the calibration process has been performed (which affect BW of the loop).

10.4 Using an FM signal generator for calibration

In principle, a signal generator can be used to improve the uncertainty of a phase noise measurement system through the use of its FM modulation capability. The signal generator can be included as the reference source, or it can be used as a temporary substitute in place of the reference source provided care is taken to ensure the levels are the same.

Signal generators typically offer accuracy specifications of 5% for the FM systems, though in practice most are considerably better. Modulation meters can measure modulation to better

than 0.5%. Unfortunately, all these measurements are taken at deviations that are not useful in phase noise measuring systems. If a signal was applied with a few kHz of deviation it would certainly overload the measurement system.

If the deviation of the signal is reduced to a useful level for calibration purposes, say 0.1 Hz deviation at 1 kHz rate (leading to a sideband level of -86 dBc measured on the test system), the accuracy is very unlikely to be 5%. A typical signal generator usually has a clause like:

$\pm 5\% \pm 1\text{Hz}$ excluding residual FM

in order to exclude residual deviations being generated by the modulation oscillator through breakthrough paths or ground currents. The generator may meet the 5% claim, but it is not guaranteed.

A safer approach is to use a modulation oscillator to drive the external FM connection of the generator. Make sure any ALC system is turned off and set the deviation to a convenient value (e.g. 1 kHz for a 1 Volt input). Set the modulation oscillator to the nominal level and then check the deviation is correct (if required use a modulation analyzer to provide a correction value). Then add a calibrated attenuator (e.g. 60 dB) between the modulation oscillator and the external FM connection of the signal generator to reduce the deviation on the RF output.

The attenuator pad can be checked using a good AC voltmeter if there are concerns about the impedance mismatches in the system. Typically, the attenuator is best operated in a 50 ohm network, which will require the external FM port of the signal generator to be externally terminated in 50 ohms, since most have a high input impedance.

Using this method, the accuracy of the FM signal from the signal generator can be accurately defined, and that information can be used to calculate the calibration values for the system. The accuracy of the sensitivity of the system (Hz/Volt or radians/Volt) can be determined to a level that can be significantly better than that achieved by the spectrum analyzer measurements.

Linearity issues associated with the mixer (phase detector) in the system will be minimized compared to the alternative ways of assessing the sensitivity.

10.5 Relative Accuracy of Systems

The author's experience of phase noise measurements has been that the most consistent and reliable measurements have been obtained using the FM Discriminator Method because it is largely unaffected by many of the factors described above. However, it is likely that methods based on digitizer methods will offer the most reliable measurements in the future.

11 Other Measurements

11.1 Two Port Phase Noise Measurements

The previous sections have all concentrated on the measurement of CW signal sources. Some of the methods can allow measurement of sources with a limited amount of FM present.

Measurements of this sort are the most common type of phase noise measurement made, and they are usually undertaken because the dominant source of noise in most systems is due to the oscillators in the system. This is particularly true close to the carrier frequency, where the oscillators in the source can be visualized as converting thermal noise to phase noise, giving a phase noise which typically rises by 6 dB or more per octave as the offset frequency drops to zero (ie as the measurement frequency approaches the carrier frequency).

There are, however, other types of phase noise measurement carried out on components and assemblies, one being two port phase noise measurements.




Figure 28: Two port additive phase noise

For two port measurements, the intention is to find the noise that a device adds to a source when the source is passed through it. The two-port device is typically an amplifier, but in principle, any device with an input and an output (including passive devices) meets the definition of a two-port device.

The most obvious noise that is added by a two-port device is the broadband noise and thermal noise. For an amplifier the thermal noise (approximately -174 dBm/Hz, in a 50 ohm system at room temperature) is added to by the noise figure of the amplifier. Noise figure is best measured by noise figure measuring methods described elsewhere in the literature and supported by test instruments substantially different to the ones described in this guide. The broadband noise added this way will appear as a mix of AM noise and phase noise (in equal contributions).

However, this is not the only noise that devices can add to a source. The noise figure measurements describe a form of additive white noise that is substantially independent of frequency offset. Other sources of noise come from fluctuations in the propagation delay of signals through a device, perhaps brought on by low frequency noise in the bias circuits for the device. The noise can also be level dependent – as an amplifier starts to enter compression, for example, it may convert amplitude noise into larger changes in phase or delay.

If it is assumed that the two port device under test does not provide frequency translation (e.g. a mixer) then the input and output to the device are the same and the usual method of measurement is to use a variation on the phase quadrature method previously described.




Figure 29: 2-port phase noise measurement

A frequency source is split into two paths, one is used to excite the DUT at the required test level and the second is used as a reference source. The reference source and the output from the DUT are connected to an RF mixer and the two signals are arranged to be in phase quadrature so that the mixer behaves as a phase detector (in the same way as with the Phase Quadrature system previously described). To obtain phase quadrature at the mixer, a variable delay line or a phase shifter needs to be included in one of the paths (usually the reference).

To a first approximation, the phase noise of the source used cancels out in the system. In the absence of a DUT, phase noise on the reference arm is the same as phase noise on the test arm, so no phase difference is detected and the output from the mixer will stay at the same voltage. It is apparent that this set up is very similar to that used in the Delay Discriminator method, and in the same way any difference in delay between the two paths will make the system behave as an FM discriminator, converting phase noise on the source to a detectable signal at the mixer about. The higher the offset frequency of the phase noise the more efficient the conversion is. In principle the degree of noise cancellation is given by the formula:

$$\text{Cancellation (dB)} = 20 \text{ Log } (2 \pi T f_m)$$

Where T is the time delay between the two arms
and f_m is the offset frequency

Typical calculations show that for most applications the degree of cancellation is high, which is fortunate, since most devices have a considerably better phase noise performance than the oscillators in the reference source. There can be other defects in the setup that do not permit this degree of cancellation.

It is normal, in these tests, to include a length of delay line that corrects for the delay time in the DUT. In most cases, it is best to use a delay line rather than an electronic phase shifter since the noise levels being tested for are low, and an electronic phase shifter usually has more phase noise addition than the DUT.

Devices which introduce a frequency translation (mixers, frequency dividers) are more problematic to measure since the above method cannot be directly used. The simplest solution is usually to test two devices at the same time, the second device replacing the direct path. The output from each path is then at the same frequency at the input to the mixer.

11.2 Calibration of Two Port Measurement

A signal generator's modulation function cannot be used to calibrate the two-port system because the modulation cancels in the system output. The modulation can, of course, provide a good indication of the time delay in each arm of the system since if the two paths are correctly matched, the modulation will not be measurable at the output of the measurement system.

Calibration is therefore usually performed by adjusting the delay line to find the peak phase detector output, and calculating the radians per volt at the output. In most cases, the mixer will be driven at a level that ensures that mixer produces a sine wave output for a linearly changing phase (time delay) with time. The sensitivity of the system in volts per radian is then equal to the peak voltage of the sine wave from the mixer.

An alternative method is to add a small (low level) calibration signal to the reference arm of the test system, using a device such as a directional coupler. Provided the signal ratio is known (or measured), and any leakage through to the DUT arm can be accounted for, this signal will principally appear on one arm of the test setup and will cause a signal to appear on the output of the system that can be measured and used to derive the calibration factors for the system.

11.3 Measurement of Phase Noise on Pulsed Signals

Not all signals that require phase noise measurements to be performed are CW sources; some sources are pulsed. In some cases, the source itself cannot be left running continuously without risking damage to the source because of thermal limitations.

Measuring noise on such sources is problematic, and at best, the systems that perform the measurement cannot measure to high levels of performance.

The first problem arises because the spectrum of the signal will have coherent signals at offsets that are multiples of the pulse repetition frequency. If, for example, the source is turned on for 10 μ s every 1 ms the signal spectrum will exhibit spurious signals at 1 kHz intervals. The amplitude of the signal is shaped by the classic $\sin(x)/x$ envelope with the first null at the pulse width of the signal (in this example 100 kHz). Any successful measurement of the phase noise will display these signals.

The system will also lose measurement sensitivity because the signal is present (in this case) for only 1 % of the time. During the off period, the system will simply reflect the residual noise floor of the system with no signal applied.

Phase locking the signal to maintain quadrature during the signal burst is also more complicated because the wanted signal is only present during part of the time. The loop can be designed to cope with these conditions, but the loop gain is severely affected by the pulsed nature of the signal.

The FM discriminator method is unsuitable because, during the off period the limiter will try to recover the signal level, and in all probability will inject noise into the measurement system while it is missing. Its behavior is unpredictable and the effects of noise bursts that result cannot be removed from the system.

For the Quadrature Method and the FM delay Line Method, measurements can be performed but it is likely that the dynamic range will be seriously reduced. In addition, video breakthrough can effect the resulting measurement by overloading the low noise amplifiers used in the system to recover the low signal levels.

Digitizer based systems can be used to make measurements by gating out the times when the signal is not present. However, it is a complex process if the signal is to be made to look continuous, and some residual artifacts are inevitable.

12 References

The following is a list of publications that the reader may find helpful as additional information to this guide.

“Phase Noise in Signal Sources” W.P. Robbins published on behalf of IEE
ISBN 0-906048-76-1

“Microwave and Wireless Synthesizers Theory and Design” Ulrich L. Rohde
ISBN 0-471-52019-5

“All you need to know about SINAD measurements using the 2023” by David Owen, available on the Aeroflex web site www.aeroflex.com

“PXImate” available from the Pickering Interfaces web site www.pickeringtest.com

Application Note 150 “Spectrum Analysis Basics” on the Agilent web site at www.agilent.com

RF Data Mate available through www.aeroflex.com

Product Note “Agilent PSA Performance Spectrum Analyzer Series Amplitude Accuracy” on the Agilent web site www.agilent.com

13 Glossary

AM. Amplitude modulation

BW. Bandwidth

FM. Frequency modulation.

EVM. Error Vector Magnitude. A measure of modulation accuracy on complex modulation systems where both the amplitude and the phase of the signal vary with time.

FM Discriminator. A device whose output voltage is dependent on the difference between the input signal frequency and a nominal centre frequency. The greater the frequency difference, the greater the output voltage. By convention, it generates a positive output voltage for input frequencies above the centre frequency and negative voltage for those below the reference.

Injection Locking. Effect when two oscillators with very similar frequencies lock their frequencies together because one sees small amounts of the other signal in its resonant circuit.

IF. Intermediate Frequency.

LNA. Low Noise Amplifier

Noise bandwidth. Equivalent bandwidth of a filter when used to measure white noise by comparing it to a filter which has perfect selectivity.

Mixer. Frequency translation device with a LO, Reference and IF output. In this guide generally assumed to be a double balanced mixer used as a down converter with a DC coupled IF output. Information on mixers can be obtained from the web site www.minicircuits.com

PCI. Peripheral Component Interconnect. Standard for PC cards. See www.pcisig.com

cPCI. Compact PCI

PXI. PCI Extension for Instruments. See www.pxisa.org

Phase Detector. A device whose output voltage is proportional to the phase difference between two input sine waves. In many cases this function is performed by a mixer in this guide.

PLL. Phase Locked Loop.

Quadrature. Refers to a condition where two signals are 90° out of phase.

Reference Source (oscillator). Source whose phase noise is generally better than the noise of the device under test.

RMS. Root Mean Square. Indicates the power contained in a signal.

SSB. Single Sideband

VCO. Voltage Controlled Oscillator

VCXO. Voltage Controlled Crystal Oscillator.

VXI. VME Extension for Instruments. \See www.vxi.org

