Phase Noise and Frequency Stability in Oscillators from RF/Microwaves to Optics

Enrico Rubiola
http://rubiola.org

Lecture series for PhD Students, Engineers, Postdoc Fellows, and Young Scientists
Guests are welcome

Webinar available on request

Purposes and General Information

This course derives from a series of seminars given at the *Tutorial Session* of international conferences of time and frequency, and from invited seminars given at INRIM (Torino, Italy), University of Pisa (Italy), NASA JPL (Pasadena, CA), DESY (Hamburg, Germany), LBNL (Berkley, CA), ANL (Chicago, IL), MPI-QO (Munich, Germany).
Some topics are of quite broader scope than oscillators, lasers, phase noise and frequency stability.
Upon request, more lectures are available, related to time and frequency and precision instrumentation, or an extended and deeper version of these lectures. Just ask if the topic you like is available.

Audience

This course is intended as a ***must*** for all the PhD students working on a topic broadly related to ***time and frequency***, including ***optics*** and ***astronomy***. Other people may be interested, on the ground of the relevance of the domain. The first lecture is of surprisingly broad interest. Then, "the appetite comes with eating," says an ancient Latin proverb.

Guests are welcome. Lectures are open to young scientists, engineers, practitioners, etc. Employees of government labs and of private companies, self-employed, hobbyists are welcome.

Prerequisites

It is understood that the attendee has a background suitable to a PhD thesis in experimental science or in engineering, or some years of experience in the domain. It is also understood that the attendee has a reasonable understanding of physics and electricity.

Language

All lectures are given in English, and all the learning material is written in English. However, the instructor understands well Italian (mother tongue) and French.

Burden and Schedule

The total burden is of 15 hour lectures, scheduled in two sessions per week, 1.5 hours each. Detailed schedule details and rooms are on the Enrico's home page http://rubiola.org, on the left-hand grey area, click on --> Syllabus --> PhD Lectures.

Learning Material

- All the slideshows (600--700 pages) are distributed via the Enrico's home page http://rubiola.org, on the left-hand grey area, click on "Syllabus" --> "PhD Lectures" (also "News," on the left-hand grey area).
- A few white papers and hundreds of copyright-free slides are available on the Enrico's home page, under "seminar slides ≥1H" and "open literature."
- Enrico Rubiola, *Phase Noise and Frequency Stability in Oscillators*, Cambridge University Press 2008 (hardbound) and 2010 (paperback). A Chinese edition is also available, 2014.
- A few chapters of the forthcoming book *The Measurement of AM and PM Noise* are available to some attendees, aware that this is still a raw draft (yet, >500 pages).

Contents

Time, and equivalently Frequency, is the most precisely measured physical quantity. Accuracy span from 10^{-5} (wrist watch) to parts in 10^{-16} (fundamental metrology). It is therefore inevitable that virtually all domains of engineering and physics rely on time-and-frequency metrology and thus need reference oscillators. We focus on the basic material, described underneath.
The clock signal and its fluctuations. Fourier statistics. The measurement of power spectra. Allan variance and other wavelet variances. AM and PM noise in devices. The Leeson effect (phase-to-frequency noise conversion in oscillators and lasers). The Pound Drever Hall frequency control (locking an oscillator or a laser to a reference cavity). Experimental methods. Time-to-digital and frequency-to-digital converters. AM and PM noise in digital systems.

Some of these topics are of broader interest than oscillators, phase noise and frequency stability. An extended summary is found in this document, starting on Page 4.

The Instructor

Enrico Rubiola is an internationally recognized scientist in the field of oscillators, frequency stability, AM-PM noise from the low RF region to optics, and precision instruments. Born in Italy in 1957, he has been a researcher with the Politecnico di Torino, a guest professor with the University of Parma, Italy, and a full professor with the University Henri Poincaré, Nancy. In 2005, he joined the University of Franche Comté and the Femto-ST Institute, and he lived in the USA in the meanwhile. He has investigated on various topics of electronics and metrology, like navigation systems, time and frequency comparisons, atomic frequency standards, and gravity.
When the French Programme d'Investissement d'Avenir (PIA) was launched in 2010, Enrico Rubiola took in charge the leadership of the local proposals related time and frequency. All the three proposals have been awarded (Oscillator IMP, First-TF, and Refimeve+).
Prof. Rubiola has authored or co-authored more than 200 articles in international journals, conferences and edited books, and has written three books. One more is in progress. He serves as a reviewer for a dozen of journals of electrical engineering, physics and optics. A wealth of articles, slides, and open literature is available on the Enrico's home page http://rubiola.org

More Details

The Measurement of Power Spectra

Basics of spectral analysis. Inside the spectrum analyzer and the FFT analyzer. Measurement time, frequency resolution, spectral leakage, background noise, precision, accuracy, etc. Mathematics joins electronics and dirty tricks. In synthesis, it took me ten years to learn it from the experiments.

Fourier Statistics, and the Cross-Spectrum Experimental Method

A physical quantity c(t) is measured with two separate instruments, each of which adds its noise. Thus, the available signals are x(t) = c(t)+a(t) and y(t) = c(t)+b(t), where a(t) and b(t) are the instrument noise. All the signals are assumed to be stationary and ergodic, which means that the physical experiment is repeatable and reproducible. By correlating and averaging the two outputs x(t) and y(t), and assuming that the two instruments are independent, it is possible to extract the statistical properties of c(t) and to reject the instrument noise. Thanks to the Wiener-Khinchin theorem, the average product of the Fourier transform of x(t) and y(t) converges to the power spectrum of c(t).

The single-channel noise is rejected proportionally to the square root of the number m of averages, and ultimately to the square root of the measurement time. Of course, the two channels must be independent. The background noise is limited by the thermal inhomogeneity of the system instead of the absolute temperature. The observation of the cross-spectrum as a function of m enables the validation of the result in some weird cases, in which a low-noise reference is not available (AM noise, laser RIN, etc.).

A major improvement results from combining the correlation methods with other experimental methods, like the bridge measurement, the differential measurement, and the synchronous detection. Applying these ideas to phase noise measurements, a background noise of parts in 10^{-21} rad^2/Hz (white) and of 10^{-18} rad^2/Hz (flicker at 1 Hz) has been reported. The latter value, turned into a length fluctuation through the wavelength of the 9.2 GHz signal, is equivalent to 4.9×10^{-12} m. This is more than 10 times smaller than the Bohr radius of the electron.

Applications and examples

The cross-spectrum method is the basis of the correlation receiver used in radio-astronomy, with which R. Hanbury-Brown measured the first radio sources in the Cassiopeia and Cygnus constellations. The correlation radiometer followed, opening the way to the re-definition of the temperature in terms of fundamental constants. Batteries and of other dc references has been measured in this way, and of course the PM and AM noise of RF/microwave signals, microwave photonic signals, and laser RIN. In semiconductor technology, small random signals reveal impurities, defects and energy traps of a dc-biased sample. Another exotic application is the measurement of electro-migration in metals at high current density, through the asymmetry between AM and PM 1/f noise, which impacts on VLSI technology.

Phase Noise and Frequency Stability

Random phase fluctuations, referred to as phase noise and closely related to frequency stability, affect precision and accuracy of timing. Random amplitude fluctuations, far less studied, may limit the most demanding experiment and systems. These types of noise impacts on numerous fields and applications, like metrology, physics, digital electronics, radars, telecommunications, optics, microwave photonics, gravitation measurements, particle accelerators, etc. Phase noise is usually described in terms of one-sided power spectral density (PSD) $S\varphi(f)$ of the random phase $\varphi(t)$. Other quantities often used and related to $S\varphi(f)$ are the PSD of the fractional frequency fluctuation y(t), denoted with Sy(f), and the two-sample (Allan) variance $\sigma^2_y(\tau)$, as a function of the measurement time τ. The same apply to the fractional amplitude fluctuation α(t).

Besides obvious thermal noise and shot noise, AM and PM noise rises from near-dc phenomena that modulate the system parameter. This describes flicker, and also the fluctuation of the environment. The rules to propagate AM-PM noise through a system depend on the noise type, and may be surprising.

PM noise is often measured by converting $\varphi(t)$ into a voltage with a mixer. The measurement of oscillators requires a reference, either another oscillator or a discriminator.

Ultimate sensitivity is achieved with the bridge (interferometric) method. After suppressing the carrier by adding an equal and opposite signal, the noise sidebands are amplified and converted to near-dc by synchronous detection.

Correlation and averaging helps rejecting the instrument noise when the signal is measured with two separate (statistically independent) instruments, each of which adds its noise. The sensitivity is limited by the thermal homogeneity, instead of the absolute temperature. This method is of special interest in the measurement of AM noise and laser RIN because even if the detector has sufficient sensitivity, we cannot validate the instrument without a reference low-noise source.

Phase Noise and Jitter in Digital Electronics

Timing analysis is generally driven by the need to ensure the consistency of logic functions. Specs like “the input must be stable 0.8 ns before the active clock edge” are just countless. To this extent, the fluctuations are analyzed with the sole purpose of making sure that priorities are respected conservatively. Time fluctuations are usually described in terms of *jitter*, which is related to the spectral measure (power spectral density) integrated over the noise bandwidth. Broadly speaking, the lower limit is determined by the longest propagation time (μs to ms); and the upper limit by the analog bandwidth, in turn related to the maximum switching frequency (GHz). In this conditions, the jitter is chiefly determined by the white noise, while flicker and other slower processes are of little or no importance.

When the design comes to the phase noise of highly stable oscillators, and to the spectral analysis of phase fluctuations, things change radically. This is the world of radars, where every pixel of the spectral measure counts because it represents a class of speed that goes in the identification of the target, and in the rejection of the clutter.

This is also the world of telecom.

We analyze the inside of components, and we give simple interpretation rules.

The Allan Variance and other Wavelet Variances

The variance, as all science students learn quite early in their university training, is a powerful tool. However, the variance cannot be used to describe diverging processes like random walk and drift. For example, the variance of the sequence 0, 0.01, 0.02, 0.03… depends on the number of terms we consider and diverges for large sequence, try yourself. Adding an appropriate weight function in the evaluation of the variance solves the problem, and results in a statistical tool suitable to the analysis of a class of divergent processes of great interest. The weight function has zero variance and bounded support, the main properties of a wavelet.

High-Resolution Time and Frequency Counters

Virtually all domains of physics and engineering at some point rely on time-and-frequency metrology, and in turn need high resolution counters.

The early instruments are based on the direct counting of the integer number of pulses in a reference time interval. This is referred to as coarse counting. The resolution associated to counting an integer number of pulses is of one cycle of the clock frequency. Ultimately, the resolution is limited by the maximum toggling frequency of the digital technology. For example, with a 100 MHz clock it is possible to measure a time interval with a resolution of 10 ns, hence a frequency with a resolution of 10^7 at 100 ms, etc. As of 2015, small FPGAs can toggle at 1 GHz clock. Higher resolution is obtained by interpolating the clock edges.

The simplest interpolator is the multi-tap delay line. Implementation in a gate array is amazingly simple, based on the idea that the internal delay of the single gates is small and predictable. A resolution up to 100 ps is expected. Other interpolators are found in the literature, namely, the analog integrator, the frequency Vernier, and the dispersive SAW filter. Commercial counters based on these methods achieve a resolution of 1--100 ps. For reference, in a coaxial cable the light travels 200-250 µm in 1 ps. Albeit the design can be tricky, the principles are simple to understand.

A trivial way to measure a frequency is to count the number of cycles --- a fractional number in the case of interpolating counters --- over a reference time τ defined by start and stop events. At a closer look, the readout is the frequency $\nu(t)$ averaged over τ with uniform weight. A frequency counter working in this way is called Π counter. This term derives from the graphical similarity of the Greek letter Π with the rectangular (uniform) weight function.

The Π counter suffers from white noise, chiefly the trigger noise. For example, a jitter of 10 ps rms at both start and stop yields a fractional-frequency resolution of 1.41×10^{-10} at $\tau = 100$ ms.

Improved resolution is obtained by averaging on highly overlapped measurements, which is equivalent to averaging the frequency $\nu(t)$ triangular weight, and for this reason the instrument is called Λ counter. Notice that the triangular weight spans on a time 2τ instead of τ. This type of measurement is equivalent to linear regression of the input frequency, based on a series of uniformly-spaced time stamps. Averaging on m overlapped measures, the resolution improves by $\sqrt{m}$. In the above example, averaging on $m=10^4$ measures yield a resolution of 1.41×10^{-12}.

Further improved resolution is achieved with a linear regression on phase data, which implements the optimum rejection of white phase noise (the dominant noise process in wide band instruments). This is equivalent to averaging the frequency $\nu(t)$ with “cap” parabolic weight, and for this reason the instrument is called Ω counter.

The Π counter is naturally suitable to the direct evaluation of the Allan variance. By contrast, feeding a stream of data measured with a Λ counter in the formula of the Allan variance, gives the modified Allan variance if the triangles are overlapped by τ and something else if the triangles are overlapped otherwise, or if there is a dead time. Feeding the readout of an Ω counter into the formula of the Allan variance results in the Parabolic Variance PVAR. This variance has advantageous properties in the detection of physical phenomena in noise.

Given a stream of measurement taken over τ, decimation enables to get new data streams averaged on 2τ, 4τ, 8τ, etc. Decimation turns out to be tricky, and interpretation mistakes are around the corner if the instrument internal mechanisms are well understood and under full control.

The Leeson Effect – Instability and Noise in Oscillators

Simply stated, an oscillator is of a loop in which a resonator sets the oscillation frequency and an amplifier compensates for the resonator loss. The oscillation amplitude is set by gain saturation, usually in the amplifier. When phase noise is introduced in the loop, the oscillator converts it to frequency noise through a process of time-domain integration. The consequence is that the oscillator phase fluctuation diverges in the long run. This phenomenon was originally referred as the "Leeson model" after a short article published by D. B. Leeson. On my side, I prefer the term "Leeson effect" in order to emphasize that it is far more general than a simple model.

The first part of this tutorial explains the phase-to-frequency conversion mechanism as a general phenomenon inherent in the feedback, following a heuristic approach based on physical insight. There follow the relationships between the noise of the internal components (sustaining amplifier, resonator, etc.) and the phase noise at the oscillator output, or equivalently the frequency stability.

The second part is the analysis of the phase noise spectra found in the data-sheet of commercial oscillators: dielectric-resonator oscillator (DRO), whispering gallery oscillator (WGO), 5-100 MHz quartz crystal oscillators, opto-electronic oscillator (OEO). The analysis gives information on the most relevant design parameters, like the quality factor Q and the driving power of the resonator, and the flicker noise of the sustaining amplifier.

The last part shows the derivation of the oscillator phase noise formulae from the elementary properties of the resonator. Interestingly, the amplitude non-linearity, necessary for the oscillation amplitude to be stable, splits the resonator relaxation time into two time constants. The approach shown in this last part is general. It applies to all oscillators, including quartz, RLC, microwave cavity, delay-line, laser, etc.

The Pound Drever Hall Frequency Control

The PDH frequency control is a milestone in radio engineering and in spectroscopy, and a smart and powerful tool available to numerous branches of experimental science.

First published in 1946 by Robert Pound, a member of the Radiation Laboratory golden team, the `Pound control' was ported to optics by John Hall (Nobel prize in 2005) and R. Drever.

The PDH control is nowadays the standard method to control a microwave oscillator or a laser oscillator to an external frequency reference, like a dielectric resonator or a Fabry-Pérot cavity.

The unique feature of the PDH scheme, which makes it superior to all competitors, is that the path length from the locked oscillator to the reference resonator cancels, so its fluctuations. Furthermore, locking relies on a null measurement of the frequency error.
Most used in general microwave and optics, this technique is the one and only which can be used at 10^{-15}…10^{-16} frequency stability level (Femto-ST ULISS oscillator, and photonic oscillators). Applications span in a wide range: metrology, optics, spectroscopy, gravitation, space/military electronics, etc.