Handbook of Frequency Stability Analysis

W.J. Riley

NIST **National Institute of Standards and Technology** • U.S. Department of Commerce

*T*he National Institute of Standards and Technology was established in 1988 by Congress to "assist industry in the development of technology ... needed to improve product quality, to modernize manufacturing processes, to ensure product reliability ... and to facilitate rapid commercialization ... of products based on new scientific discoveries." NIST, originally founded as the National Institute of Standards in 1901, works to strengthen U.S. industry's competitiveness; advance science and engineering; and improve public health, safety, and the environment. One of the agency's basic functions is to develop, maintain, and retain custody of the national standards of measurement, and provide the means and methods for comparing standards used in science, engineering, manufacturing, commerce, industry, and education with the standards adopted or recognized by the Federal Government.

As an agency of the U.S. Commerce Department's Technology Administration, NIST conducts basic and applied research in the physical sciences and engineering, and develops measurement techniques, test methods, standards, and related services. The Institute does generic and precompetitive work on new and advanced technologies. NIST's research facilities are located at Gaithersburg, MD 20899, and at Boulder, CO 80305. Major technical operating units and their principal activities are listed below. For more information visit the NIST Website at http://www.nist.gov, or contact the Publications and Program Inquiries Desk, 301-975-3058.

Office of the Director
·National Quality Program
·International and Academic Affairs

Technology Services
·Standards Services
·Technology Partnerships
·Measurement Services
·Information Services
·Weights and Measures

Advanced Technology Program
·Economic Assessment
·Information Technology and Applications
·Chemistry and Life Sciences
·Electronics and Photonics Technology

Manufacturing Extension Partnership Program
·Regional Programs
·National Programs
·Program Development

Electronics and Electrical Engineering Laboratory
·Microelectronics
·Law Enforcement Standards
·Electricity
·Semiconductor Electronics
·Radio-Frequency Technology[1]
·Electromagnetic Technology[1]
·Optoelectronics[1]
·Magnetic Technology[1]

Materials Science and Engineering Laboratory
·Intelligent Processing of Materials
·Ceramics
·Materials Reliability[1]
·Polymers
·Metallurgy
·NIST Center for Neutron Research

Chemical Science and Technology Laboratory
·Biotechnology
·Process Measurements
·Surface and Microanalysis Science
·Physical and Chemical Properties[2]
·Analytical Chemistry

Physics Laboratory
·Electron and Optical Physics
·Atomic Physics
·Optical Technology
·Ionizing Radiation
·Time and Frequency[1]
·Quantum Physics[1]

Manufacturing Engineering Laboratory
·Precision Engineering
·Manufacturing Metrology
·Intelligent Systems
·Fabrication Technology
·Manufacturing Systems Integration

Building and Fire Research Laboratory
·Applied Economics
·Materials and Construction Research
·Building Environment
·Fire Research

Information Technology Laboratory
·Mathematical and Computational Sciences[2]
·Advanced Network Technologies
·Computer Security
·Information Access
·Convergent Information Systems
·Information Services and Computing
·Software Diagnostics and Conformance Testing
·Statistical Engineering

[1]At Boulder, CO 80305
[2]Some elements at Boulder, CO

Handbook of Frequency Stability Analysis

W.J. Riley
Under contract with:

Time and Frequency Division
Physics Laboratory
National Institute of Standards and Technology
325 Broadway
Boulder, CO 80305

Dedication

This handbook is dedicated to the memory of Dr. James A. Barnes (1933−2002), a pioneer in the statistics of frequency standards.

James A. Barnes was born in 1933 in Denver, Colorado. He received a Bachelor's degree in engineering physics from the University of Colorado, a Masters degree from Stanford University, and in 1966 a Ph.D. in physics from the University of Colorado. Jim also received an MBA from the University of Denver.

After graduating from Stanford, Jim joined the National Institute of Standards, now the National Institute of Standards and Technology (NIST). Jim was the first Chief of the Time and Frequency Division when it was created in 1967 and set the direction for this division in his 15 years of leadership. During his tenure at NIST Jim made many significant contributions to the development of statistical tools for clocks and frequency standards. Also, three primary frequency standards (NBS 4, 5, and 6) were developed under his leadership. While he was division chief, closed-captioning was developed (which received an Emmy award) and the speed of light was measured. Jim received the NBS Silver Medal in 1965 and the Gold Medal in 1975. In 1992, Jim received the I.I. Rabi Award from the IEEE Frequency Control Symposium "for contributions and leadership in the development of the statistical theory, simulation and practical understanding of clock noise, and the application of this understanding to the characterization of precision oscillators and atomic clocks." In 1995, he received the Distinguished PTTI Service Award. Jim was a Fellow of the IEEE. After retiring from NIST in 1982, Jim worked for Austron.

Jim Barnes died Sunday, January 13, 2002, in Boulder, Colorado after a long struggle with Parkinson's disease. He was survived by a brother, three children, and two grandchildren.

Note: This biography is published with permission and taken from his memoriam on the UFFC web site at: http://www.ieee-uffc.org/fcmain.asp?page=barnes.

Preface

I have had the great privilege of working in the time and frequency field over the span of my career. I have seen atomic frequency standards shrink from racks of equipment to chip scale, and be manufactured by the tens of thousands, while primary standards and the time dissemination networks that support them have improved by several orders of magnitude. During the same period, significant advances have been made in our ability to measure and analyze the performance of those devices. This Handbook summarizes the techniques of frequency stability analysis, bringing together material that I hope will be useful to the scientists and engineers working in this field.

_____

I acknowledge the contributions of many colleagues in the Time and Frequency community who have contributed the analytical tools that are so vital to this field. In particular, I wish to recognize the seminal work of J.A. Barnes and D.W. Allan in establishing the fundamentals at NBS, and D.A. Howe in carrying on that tradition today at NIST. Together with such people as M.A. Weiss and C.A. Greenhall, the techniques of frequency stability analysis have advanced greatly during the last 45 years, supporting the orders-of-magnitude progress made on frequency standards and time dissemination.

I especially thank David Howe and the other members of the NIST Time and Frequency Division for their support, encouragement, and review of this Handbook.

Contents

1   Introduction

This handbook describes practical techniques for frequency stability analysis. It covers the definitions of frequency stability, measuring systems and data formats, pre-processing steps, analysis tools and methods, post-processing steps, and reporting suggestions. Examples are included for many of these techniques. Some of the examples use the Stable32 program [1], which is a tool for studying and performing frequency stability analyses. Two general references [2,3] for this subject are also given.

This handbook can be used both as a tutorial and as a reference. If this is your first exposure to this field, you may find it helpful to scan the sections to gain some perspective regarding frequency stability analysis. I strongly recommend consulting the references as part of your study of this subject matter. The emphasis is on time domain stability analysis, where specialized statistical variances have been developed to characterize clock noise as a function of averaging time. I present methods to perform those calculations, identify noise types, and determine confidence limits. It is often important to separate deterministic factors such as aging and environmental sensitivity from the stochastic noise processes. One must always be aware of the possibility of outliers and other measurement problems that can contaminate the data.

Suggested analysis procedures are recommended to gather data, preprocess it, analyze stability, and report results. Throughout these analyses, it is worthwhile to remember R.W. Hamming's axiom that "the purpose of computing is insight, not numbers." Analysts should feel free to use their intuition and experiment with different methods that can provide a deeper understanding.

References for Introduction

1.   The Stable32 Program for Frequency Stability Analysis, Hamilton Technical Services, Beaufort, SC 29907, http://www.wriley.com.
2.   D.B Sullivan, D.W Allan, D.A Howe, and F.L Walls, eds., "Characterization of clocks and oscillators," Natl. Inst. Stand. Technol. Technical Note 1337,
      http://tf.nist.gov/timefreq/general/pdf/868.pdf (March 1990).
3.   D.A. Howe, D.W. Allan, and J.A. Barnes, "Properties of signal sources and measurement methods," *Proc. 35th Freq. Cont. Symp*., pp. 1-47,
      http://tf.nist.gov/timefreq/general/pdf/554.pdf (May 1981).

2   Frequency Stability Analysis

The time domain stability analysis of a frequency source is concerned with characterizing the variables x(t) and y(t), the phase (expressed in units of time error) and the fractional frequency, respectively. It is accomplished with an array of phase arrays x_i and frequency data arrays y_i, where the index i refers to data points equally spaced in time. The x_i values have units of time in seconds, and the y_i values are (dimensionless) fractional frequency, $\Delta f/f$. The x(t) time fluctuations are related to the phase fluctuations by $\phi(t) = x(t) \cdot 2\pi\nu_0$, where ν_0 is the nominal carrier frequency in hertz. Both are commonly called "phase" to distinguish them from the independent time variable, t. The data sampling or measurement interval, τ_0, has units of seconds. The analysis interval or period, loosely called "averaging time", τ, may be a multiple of τ_0 ($\tau = m\tau_0$, where m is the averaging factor).

> The objective of a frequency stability analysis is to characterize the phase and frequency fluctuations of a frequency source in the time and frequency domains.

The goal of a time domain stability analysis is a concise, yet complete, quantitative and standardized description of the phase and frequency of the source, including their nominal values, the fluctuations of those values, and their dependence on time and environmental conditions.

A frequency stability analysis is normally performed on a single device, not on a population of devices. The output of the device is generally assumed to exist indefinitely before and after the particular data set was measured, which is the (finite) population under analysis. A stability analysis may be concerned with both the stochastic (noise) and deterministic (systematic) properties of the device under test. It is also generally assumed that the stochastic characteristics of the device are constant (both stationary over time and ergodic over their population). The analysis may show that this is not true, in which case the data record may have to be partitioned to obtain meaningful results. It is often best to characterize and remove deterministic factors (e.g., frequency drift and temperature sensitivity) before analyzing the noise. Environmental effects are often best handled by eliminating them from the test conditions. It is also assumed that the frequency reference instability and instrumental effects are either negligible or removed from the data. A common problem for time domain frequency stability analysis is to produce results at the longest possible analysis interval in order to minimize test time and cost. Computation time is generally not as much a factor.

2.1   Background

The field of modern frequency stability analysis began in the mid 1960's with the emergence of improved analytical and measurement techniques. In particular, new statistics became available that were better suited for common clock noises than the classic N-sample variance, and better methods were developed for high resolution measurements (e.g., heterodyne period measurements with electronic counters, and low noise phase noise measurements with double-balanced diode mixers). A seminal conference on short-term stability in 1964 [1], and the introduction of the two-sample (Allan) variance in 1966 [2] marked the beginning of this new era, which was summarized in a special issue of the *Proceedings of the IEEE* in 1966 [3]. This period also marked the introduction of commercial atomic frequency standards, increased emphasis on low phase noise, and the use of the LORAN radio navigation system for global precise time and frequency transfer. The subsequent advances in the performance of frequency sources depended largely on the improved ability to measure and analyze their stability. These advances also mean that the field of frequency stability analysis has become more complex. It is the goal of this handbook to help the analyst deal with this complexity.

An example of the progress that has been made in frequency stability analysis from the original Allan variance in 1966 through Thêo1 in 2003 is shown in the plots below. The error bars show the improvement in statistical confidence for the same data set, while the extension to longer averaging time provides better long-term clock characterization without the time and expense of a longer data record.
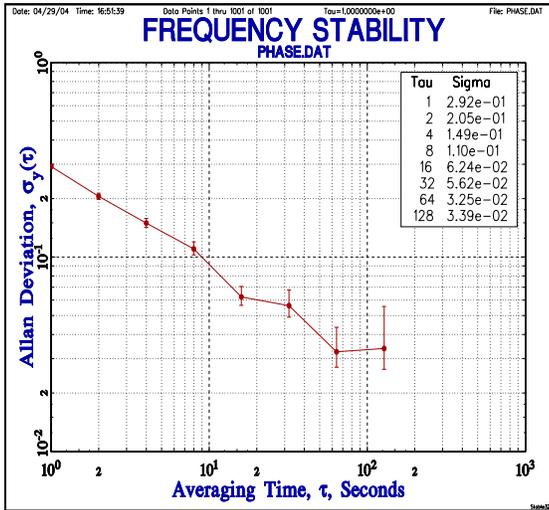
Original Allan (a)
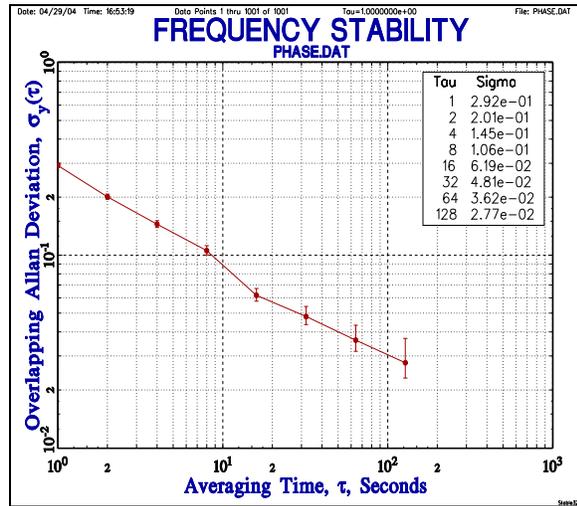
Overlapping Allan (b)
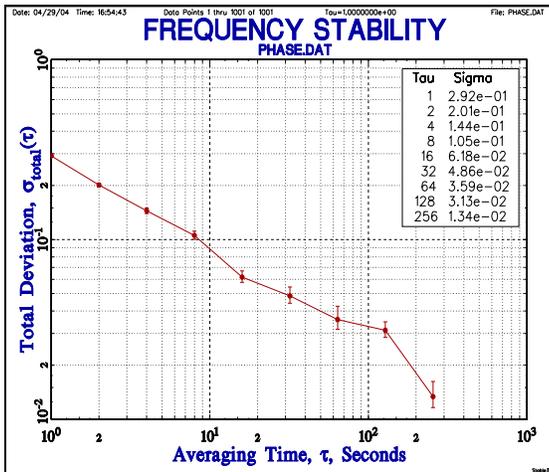
Total (c)
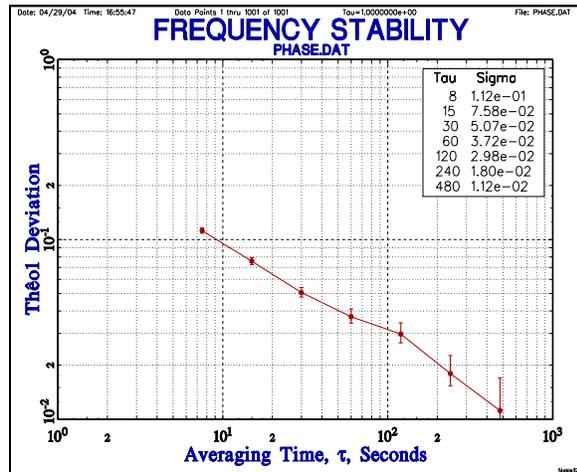
Thêo1 (d)
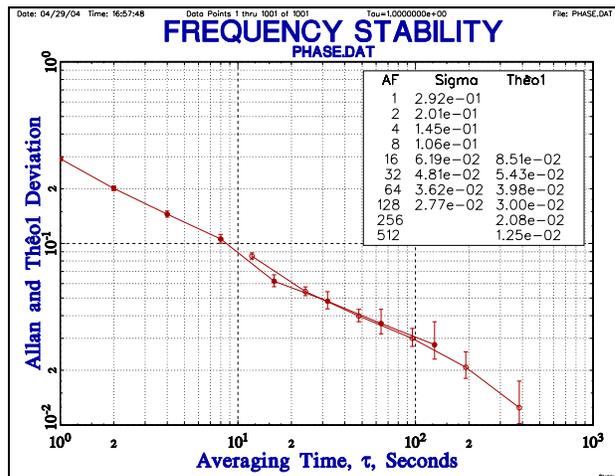
Overlapping and Thêo1 (e)

Figure 1. Progress in frequency stability analysis. (a) Original Allan. (b) Overlapping Allan. (c) Total. (d) Thêo1. (e) Overlapping and Thêo1.

This handbook includes detailed information about these (and other) stability measures.

References for Frequency Stability Analysis

1. Proc. of the IEEE-NASA Symposium on the Definition and Measurement of Short-Term Frequency Stability, *NASA SP-80*, (Nov. 1964).
2. D.W. Allan, "The Statistics of Atomic Frequency Standards," *Proc. IEEE*, 54(2): 221-230(Feb. 1966).
3. Special Issue on Frequency Stability, *Proc. IEEE*, 54(2)(Feb. 1966).

3   Definitions and Terminology

The field of frequency stability analysis, like most others, has its own specialized definitions and terminology. The basis of a time domain stability analysis is an array of equally spaced phase (really time error) or fractional frequency deviation data arrays, x_i and y_i, respectively, where the index i refers to data points in time. These data are equivalent, and conversions between them are possible. The x values have units of time in seconds, and the y values are (dimensionless) fractional frequency, $\Delta f/f$. The $x(t)$ time fluctuations are related to the phase fluctuations by $\phi(t) = x(t) \cdot 2\pi v_0$, where v_0 is the carrier frequency in hertz. Both are commonly called "phase" to distinguish them from the independent time variable, t. The data sampling or measurement interval, τ_0, has units of seconds. The analysis or averaging time, τ, may be a multiple of τ_0 ($\tau = m\tau_0$, where m is the averaging factor). Phase noise is fundamental to a frequency stability analysis, and the type and magnitude of the noise, along with other factors such as aging and environmental sensitivity, determine the stability of the frequency source.

> Specialized definitions and terminology are used for frequency stability analysis.

3.1.   Noise Model

A frequency source has a sine wave output signal given by [1]

$$V(t) = [V_0 + \varepsilon(t)]\sin[2\pi v_0 t + \phi(t)], \tag{1}$$

where V_0 = nominal peak output voltage
        $\varepsilon(t)$ = amplitude deviation
        v_0 = nominal frequency
        $\phi(t)$ = phase deviation.

For the analysis of frequency stability, we are concerned primarily with the $\phi(t)$ term. The instantaneous frequency is the derivative of the total phase:

$$v(t) = v_0 + \frac{1}{2\pi}\frac{d\phi}{dt}. \tag{2}$$

For precision oscillators, we define the fractional frequency as

$$y(t) = \frac{\Delta f}{f} = \frac{v(t) - v_0}{v_0} = \frac{1}{2\pi v_0}\frac{d\phi}{dt} = \frac{dx}{dt}, \tag{3}$$

where

$$x(t) = \phi(t)/2\pi v_0. \tag{4}$$

3.2.   Power Law Noise

It has been found that the instability of most frequency sources can be modeled by a combination of power-law noises having a spectral density of their fractional frequency fluctuations of the form $S_y(f) \propto f^{\alpha}$, where f is the Fourier or sideband frequency in hertz, and α is the power law exponent.

Noise Type	**α**
White PM (W PM)	2
Flicker PM (F PM)	1
White FM (W FM)	0
Flicker FM (F FM)	−1
Random Walk FM (RW FM)	−2

Flicker Walk FM (FW FM)   −3
Random Run FM (RR FM)   −4

Examples of the four most common of these noises are shown in Table 1.

Table 1. Examples of the four most common noise types.
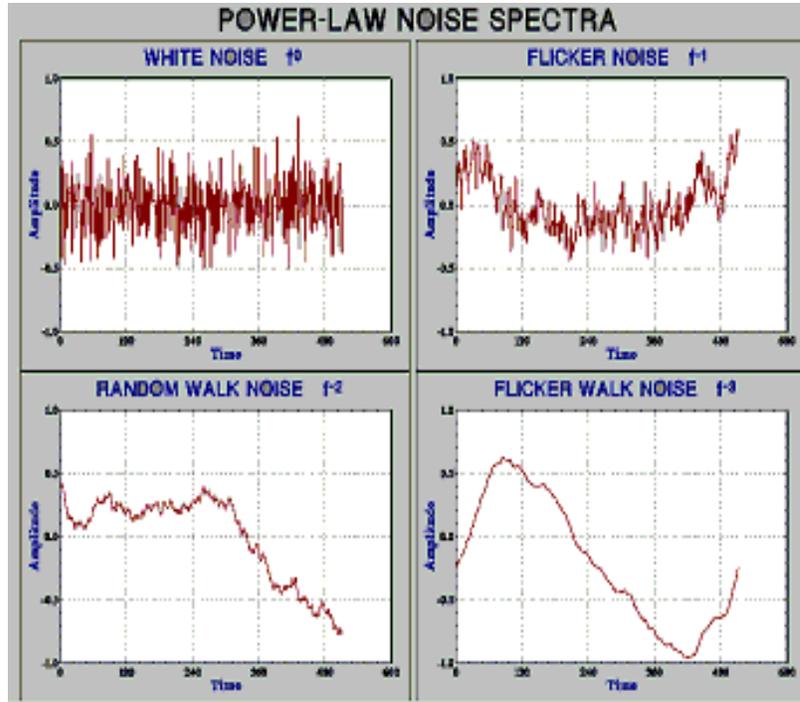
3.3.   Stability Measures

The standard measures for frequency stability in the time and frequency domains are the overlapped Allan deviation, $\sigma_y(\tau)$, and the SSB phase noise, $£(f)$, as described in more detail later in this handbook.

3.4.   Differenced and Integrated Noise

Taking the differences between adjacent data points plays an important role in frequency stability analysis for performing phase to frequency data conversion, calculating Allan (and related) variances, and doing noise identification using the lag 1 autocorrelation method [2]. Phase data $x(t)$ may be converted to fractional frequency data $y(t)$ by taking the first differences $x_{i+1} - x_i$ of the phase data and dividing by the sampling interval τ. The Allan variance is based on the first differences $y_{i+1} - y_i$ of the fractional frequency data or, equivalently, the second differences $y_{i+2} - 2y_{i+1} + y_i$ of the phase data. Similarly, the Hadamard variance is based on third differences $x_{i+3} - 3x_{i+2} + 3x_{i+1} - x_i$ of the phase data.

Taking the first differences of a data set has the effect of making it less divergent. In terms of its spectral density, the α value is increased by 2. For example, flicker FM data ($\alpha = -1$) is changed into flicker PM data ($\alpha = +1$). That is the reason that the Hadamard variance is able to handle more divergent noise types ($\alpha \geq -4$) than the Allan variance ($\alpha \geq -2$) can. It is also the basis of the lag 1 autocorrelation noise identification method whereby first differences are taken until α becomes ≥ 0.5. The plots below show random run noise differenced first to random walk noise and again to white noise.

6

(a) Original random run (RR)
noise

(b) Differenced RR noise =
random walk (RW) noise

(c) Differenced RW noise =
white (W) noise
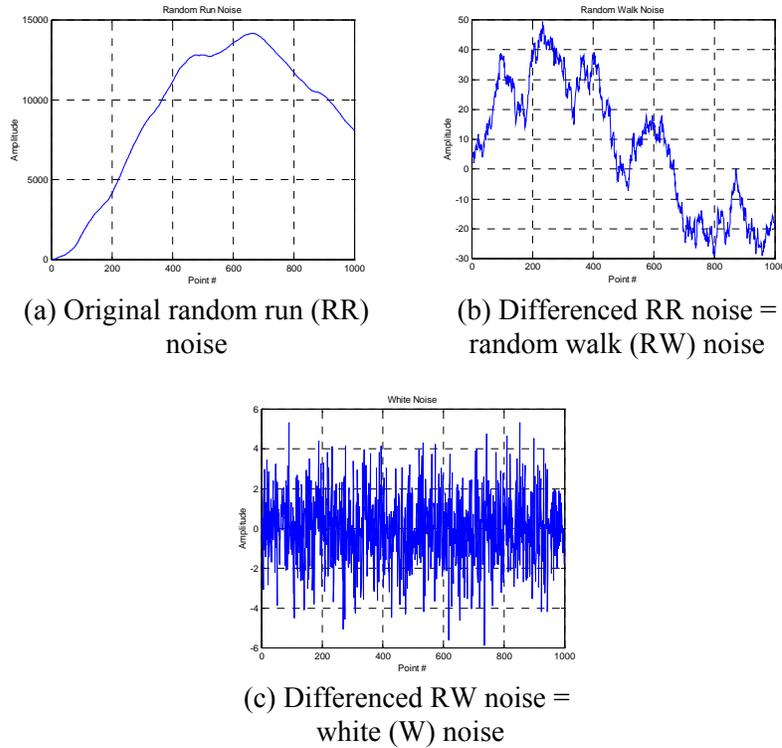
**Figure 2. (a) Random run noise, difference to (b) random walk noise
and (c) white noise.**

The more divergent noise types are sometimes referred to by their color. White noise has a flat spectral density (by analogy to white light). Flicker noise has an f^{-1} spectral density, and is called pink or red (more energy toward lower frequencies). Continuing the analogy, f^{-2} (random walk) noise is called brown, and f^{-3} (flicker walk) noise is called black, although that terminology is seldom used in the field of frequency stability analysis.

Integration is the inverse operation of differencing. Numerically integrating frequency data converts it into phase data (with an arbitrary initial value). Such integration subtracts 2 from the original α value. For example, the random run data in Figure 2(a) was generated by simulating random walk FM data and converting it to phase data by numerical integration.

3.5.   Glossary

See the Glossary chapter at the end of this handbook for brief definitions of many of the important terms used in the field of frequency stability analysis.

References for Definitions and Terminology

1.   "IEEE Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology−Random instabilities," IEEE Std. 1139 (July 1999).
2.   W.J. Riley and C.A. Greenhall, "Power law noise identification using the lag 1 autocorrelation," *Proc. 18th European Frequency and Time Forum*, University of Surrey, Guildford, U.K. (April 5−7, 2004).

4   Standards

Standards have been adopted for the measurement and characterization of frequency stability, as shown in the references below [1-5]. These standards define terminology, measurement methods, means for characterization and specification, etc. In particular, IEEE-Std-1139 contains definitions, recommendations, and examples for the characterization of frequency stability.

Several standards apply to the field of frequency stability analysis.

References for Standards

1. "Characterization of frequency and phase noise, Intl. Consult. Comm. (C.C.I.R.), Report 580," pp. 142-150 (1986).
2. MIL-PRF-55310, "Oscillator, crystal controlled, general specification for (2006)."
3. R.L. Sydnor, ed., "The selection and use of precise frequency systems," ITU-R Handbook (1995).
4. "Guide to the expression of uncertainty in measurement," Intl. Stand. Org. (ISO), ISBN 92-67-10188-9 (1995).
5. "IEEE standard definitions of physical quantities for fundamental frequency and time metrology−Random instabilities," IEEE Std. 1139 (July 1999).

5   Time Domain Stability

The stability of a frequency source in the time domain is based on the statistics of its phase or frequency fluctuations as a function of time, a form of time series analysis [1]. This analysis generally uses some type of variance, a second moment measure of the fluctuations. For many divergent noise types commonly associated with frequency sources, the standard variance, which is based on the variations around the average value, is not convergent, and other variances have been developed that provide a better characterization of such devices. A key aspect of such a characterization is the dependence of the variance on the averaging time used to make the measurement, which dependence shows the properties of the noise.

> Time domain stability measures are based on the statistics of the phase or frequency fluctuations as a function of time.

5.1.   Sigma-Tau Plots

The most common way to express the time domain stability of a frequency source is by means of a sigma-tau plot that shows some measure of frequency stability versus the time over which the frequency is averaged. Log sigma versus log tau plots show the dependence of stability on averaging time, and show both the stability value and the type of noise. The power law noises have particular slopes, μ, as shown on the following log s versus log τ plots, and α and μ are related as shown in the table below:

Noise	α	μ
W PM	2	−2
F PM	1	$\sim$−2
W FM	0	−1
F FM	−1	0
RW FM	−2	1

The log σ versus log τ slopes are the same for the two PM noise types, but are different on a Mod sigma plot, which is often used to distinguish between them.

Sigma Tau Diagram

White PM
or
Flicker PM

$\sigma_y(\tau) \sim \tau^{\mu/2}$
$S_y(f) \sim f^{\alpha}$
$\mu = -\alpha-1$

Freq Drift

τ^{-1}

White FM   Flicker FM   RW FM τ^{+1}

$\tau^{-1/2}$ τ^{0} $\tau^{+1/2}$

log $\sigma_y(\tau)$

-9
-11
-13
-15

log τ
0   2   4   6   8   10

Mod Sigma Tau Diagram

White PM

Mod $\sigma_y(\tau) \sim \tau^{\mu'/2}$
$S_y(f) \sim f^{\alpha}$
$\mu' = -\alpha-1$

Flicker PM

Freq Drift

$\tau^{-3/2}$

τ^{-1}

White FM   Flicker FM   RW FM τ^{+1}

$\tau^{-1/2}$ τ^{0} $\tau^{+1/2}$

log Mod $\sigma_y(\tau)$

-9
-11
-13
-15

log τ
0   2   4   6   8
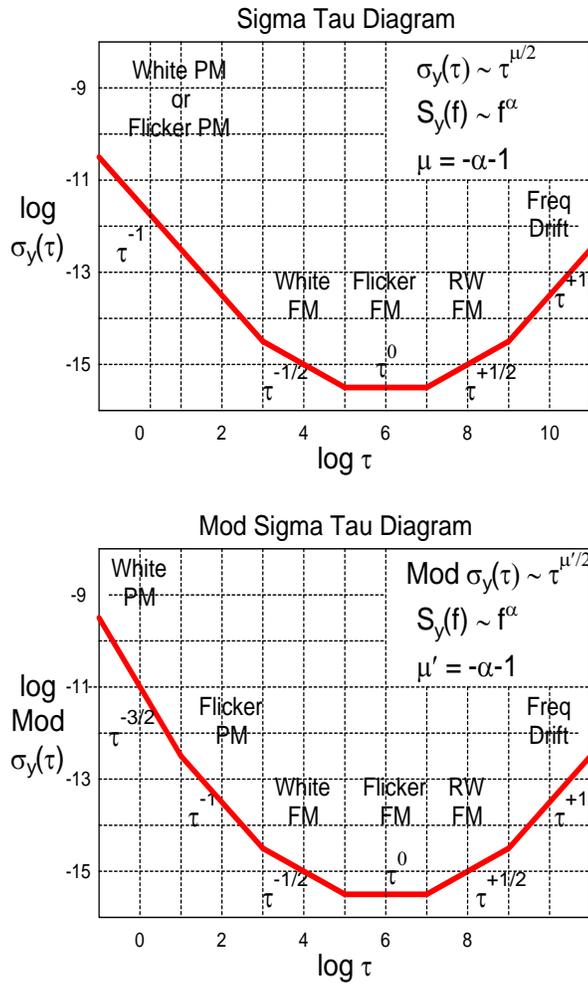
Figure 3. (a) Sigma tau diagram. (b) Mod sigma tau diagram.

5.2   Variances

Variances are used to characterize the fluctuations of a frequency source [2-3]. These are second-moment measures of scatter, much as the standard variance is used to quantify the variations in, say, the length of rods around a nominal value. The variations from the mean are squared, summed, and divided by one less than the number of measurements; this number is called the "degrees of freedom."

Several statistical variances are available to the frequency stability analyst, and this section provides an overview of them, with more details to follow. The Allan variance is the most common time domain measure of frequency stability, and there are several versions of it that provide better statistical confidence, can distinguish between white and flicker phase noise, and can describe time stability. The Hadamard variance can better handle frequency drift and more divergence noise types, and several versions of it are also available. The newer Total and Thêo1 variances can provide better confidence at longer averaging factors.

There are two categories of stability variances: unmodified variances, which use d[th] differences of phase samples, and modified variances, which use d[th] differences of averaged phase samples. The Allan variances correspond to d = 2, and the Hadamard variances to d = 3. The corresponding variances are defined as a scaling factor times the expected value of the differences squared. One obtains unbiased estimates of this variance from available phase data by computing time averages of the differences squared. The usual choices for the increment between estimates (the time

step) are the sample period τ_0 and the analysis period τ, a multiple of τ_0. These give respectively the overlapped estimator and non-overlapped estimators of the stability.

Variance Type	Characteristics
Standard	Non-convergent for some clock noises – don't use
Allan	Classic – use only if required – relatively poor confidence
Overlapping Allan	General purpose - most widely used – first choice
Modified Allan	Used to distinguish W and F PM
Time	Based on modified Allan variance
Hadamard	Rejects frequency drift, and handles divergent noise
Overlapping Hadamard	Better confidence than normal Hadamard
Total	Better confidence at long averages for Allan
Modified Total	Better confidence at long averages for modified Allan
Time Total	Better confidence at long averages for time
Hadamard Total	Better confidence at long averages for Hadamard
Thêo1	Provides information over nearly full record length
ThêoH	Hybrid of Allan and ThêoBR (bias-removed Thêo1) variances

- All are second moment measures of dispersion – scatter or instability of frequency from central value.
- All are usually expressed as deviations.
- All are normalized to standard variance for white FM noise.
- All except standard variance converge for common clock noises.
- Modified types have additional phase averaging that can distinguish W and F PM noises.
- Time variances based on modified types.
- Hadamard types also converge for FW and RR FM noise.
- Overlapping types provide better confidence than classic Allan variance.
- Total types provide better confidence than corresponding overlapping types.
- ThêoH (hybrid-ThêoBR) and Thêo1 (Theoretical Variance #1) provide stability data out to 75 % of record length.
- Some are quite computationally intensive, especially if results are wanted at all (or many) analysis intervals (averaging times), τ. Use octave or decade τ intervals.

The modified Allan deviation (MDEV) can be used to distinguish between white and flicker PM noise. For example, the W and F PM noise slopes are both ≈ -1.0 on the Allan Deviation (ADEV) plots in Figure 4, but they can be distinguished as -1.5 and -1.0, respectively, on the MDEV plots.
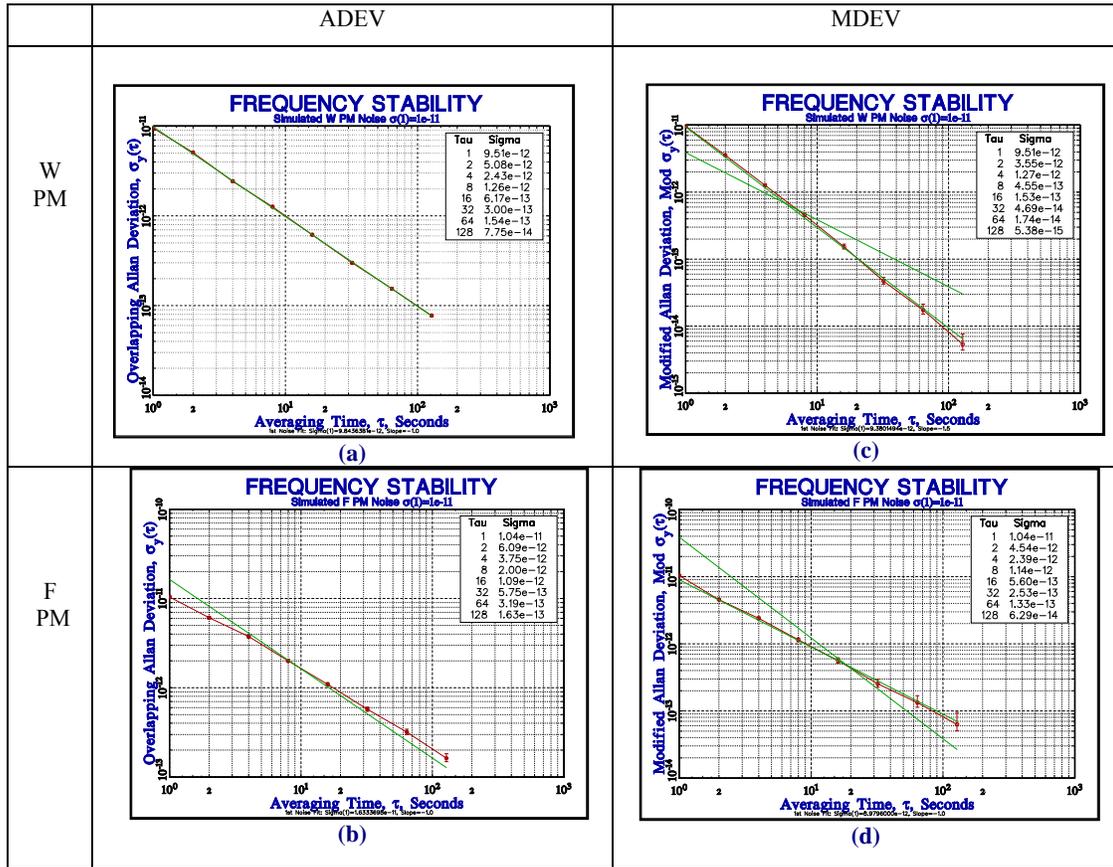
Figure 4. (a) Slope of W PM using Adev, (b) slope of F PM using ADEV, (c) slope of W PM using MDEV, and (d) slope of F PM using MDEV.

The Hadamard deviation may be used to reject linear frequency drift when a stability analysis is performed. For example, the simulated frequency data for a rubidium frequency standard in Figure 5(a) shows significant drift. Allan deviation plots for these data are shown in Figure 5(c) and (d) for the original and drift-removed data. Notice that, without drift removal, the Allan deviation plot has a +τ dependence at long τ, a sign of linear frequency drift. However, as seen in Figure 5(b), the Hadamard deviation for the original data is nearly the same as the Allan deviation after drift removal, but it has lower confidence for a given τ.

12
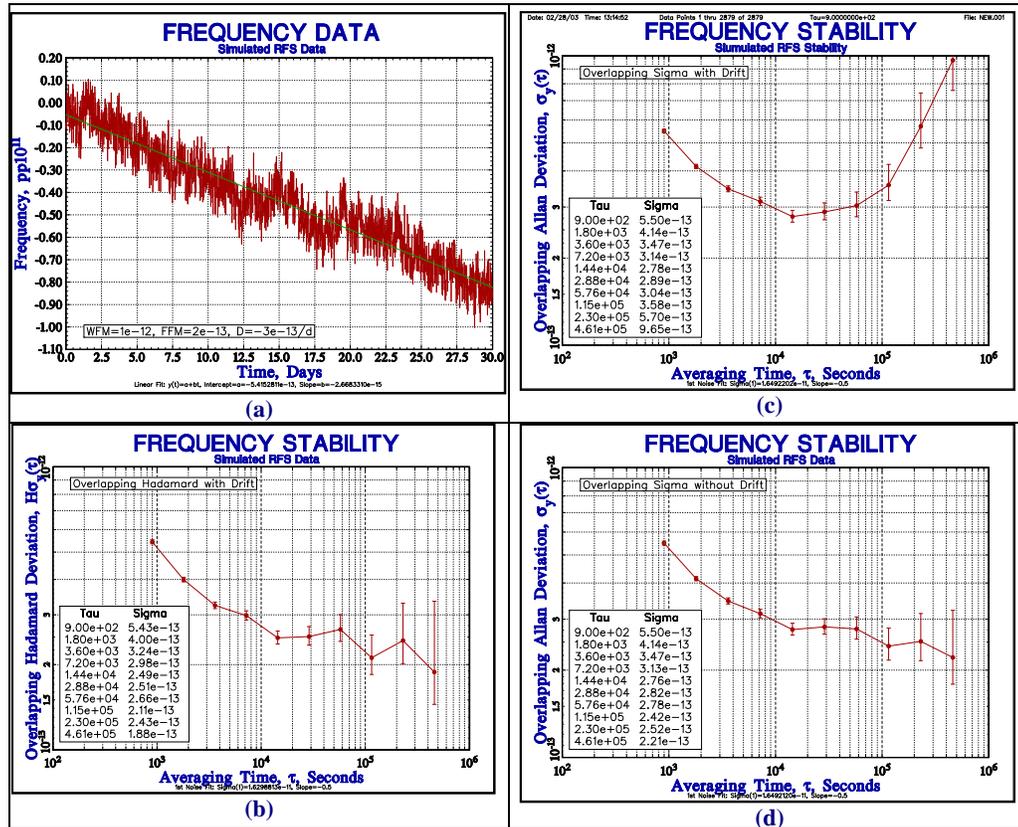
Figure 5. (a) Simulated frequency data for a rubidium frequency standard, (b) overlapping Hadamard with drift, (c) overlapping sigma with drift, and (d) overlapping sigma without drift.

References for Time Domain Stability

1. G.E.P. Box and G.M. Jenkins, <u>Time Series Analysis: Forecasting and Control</u>, Holden-Day, San Francisco (1970).
2. J. Rutman, "Characterization of phase and frequency instabilities in precision frequency sources: Fifteen years of progress," *Proc. IEEE*, 66(9): 1048-1075 (1978)
3. S.R. Stein, "Frequency and time: Their measurement and characterization," Precision Frequency Control, 2, E.A. Gerber and A. Ballato, eds., Academic Press, New York, ISBN 0-12-280602-6 (1985).

5.2.1. Standard Variance

The classic N-sample or standard variance is defined as

> The standard variance should not be used for the analysis of frequency stability.

$$s^2 = \frac{1}{N-1}\sum_{i=1}^{N}\left(y_i - \bar{y}\right)^2 , \qquad (5)$$

where the y_i are the N fractional frequency values, and $\bar{y} = \frac{1}{N}\sum_{i=1}^{N} y_i$ is the average frequency. The standard variance is usually expressed as its square root, the standard deviation, s. It is not recommended as a measure of frequency stability because it is non- convergent for some types of noise commonly found in frequency sources, as shown in the figure below.
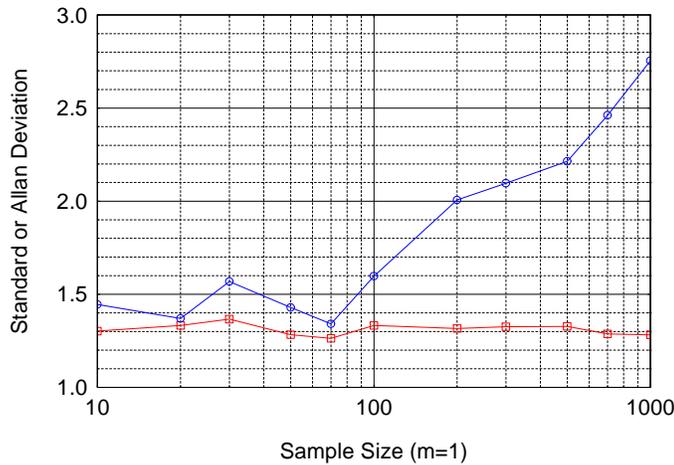
Figure 6. Convergence of standard and Allan deviation for FM noise.

The standard deviation (upper curve) increases with the number of samples of flicker FM noise used to determine it, while the Allan deviation (lower curve and discussed below) is essentially constant.

The problem with the standard variance stems from its use of the deviations from the average, which is not stationary for the more divergence noise types. That problem can be solved by instead using the first differences of the fractional frequency values (the second differences of the phase), as described for the Allan variance in Section 5.2.2.

In the context of frequency stability analysis, the standard variance is used primarily in the calculation of the B1 ratio for noise recognition.

Reference for Standard Variance

1. D.W. Allan, "Should the Classical Variance be used as a Basic Measure in Standards Metrology?" *IEEE Trans. Instrum. Meas.*, IM-36: 646-654 (1987)

5.2.2.   Allan Variance

The Allan variance is the most common time domain measure of frequency stability. Similar to the standard variance, it is a measure of the fractional frequency fluctuations, but has the advantage of being convergent for most types of clock noise. There are several versions of the Allan variance that provide better statistical confidence, can distinguish between white and flicker phase noise, and can describe time stability.

The original non-overlapped Allan, or two-sample variance, AVAR, is the standard time domain measure of frequency stability [1, 2]. It is defined as
It is defined as

> The original Allan variance has been largely superseded by its overlapping version.

$$\sigma_y^2(\tau) = \frac{1}{2(M-1)} \sum_{i=1}^{M-1} [y_{i+1} - y_i]^2 , \tag{6}$$

where y_i is the ith of M fractional frequency values averaged over the measurement (sampling) interval, τ. Note that these y symbols are sometimes shown with a bar over them to denote the averaging.

In terms of phase data, the Allan variance may be calculated as

$$\sigma_y^2(\tau) = \frac{1}{2(N-2)\tau^2} \sum_{i=1}^{N-2} [x_{i+2} - 2x_{i+1} + x_i]^2 , \qquad\qquad (7)$$

where x_i is the ith of the $N = M+1$ phase values spaced by the measurement interval τ.

The result is usually expressed as the square root, $\sigma_y(\tau)$, the Allan deviation, ADEV. The Allan variance is the same as the ordinary variance for white FM noise, but has the advantage, for more divergent noise types such as flicker noise, of converging to a value that is independent on the number of samples. The confidence interval of an Allan deviation estimate is also dependent on the noise type, but is often estimated as $\pm\sigma_y(\tau)/\sqrt{N}$.

5.2.2.   Overlapping Samples

Some stability calculations can utilize (fully) overlapping samples, whereby the calculation is performed by utilizing all possible combinations of the data set, as shown in the diagram and formulae below. The use of overlapping samples improves the confidence of the resulting stability estimate, but at the expense of

> Overlapping samples are used to improve the confidence of a stability estimate.

greater computational time. The overlapping samples are not completely independent, but do increase the effective number of degrees of freedom. The choice of overlapping samples applies to the Allan and Hadamard variances. Other variances (e.g., total) always use them.

Overlapping samples don't apply at the basic measurement interval, which should be as short as practical to support a large number of overlaps at longer averaging times.
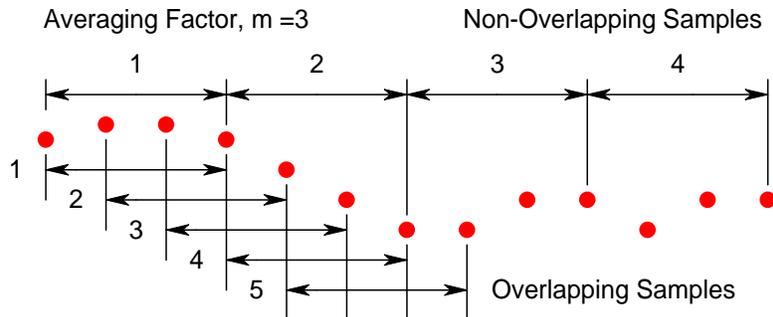
Non-Overlapped Allan Variance: Stride = τ = averaging period = $m\cdot\tau_0$

$$\sigma_y^2(\tau) = \frac{1}{2(M-1)} \sum_{i=1}^{M-1} (y_{i+1} - y_i)^2 \qquad\qquad (8)$$

Overlapped Allan Variance: Stride = τ_0 = sample period

$$\sigma_y^2(\tau) = \frac{1}{2m^2(M-2m+1)} \sum_{j=1}^{M-2m+1} \sum_{i=j}^{j+m-1} (y_{i+m} - y_i)^2 \qquad\qquad (9)$$

Figure 7. Comparison of non-overlapping and overlapping sampling.

The following plots show the significant reduction in variability, hence increased statistical confidence, obtained by using overlapping samples in the calculation of the Hadamard deviation.

Non-Overlapping Samples
Overlapping Samples
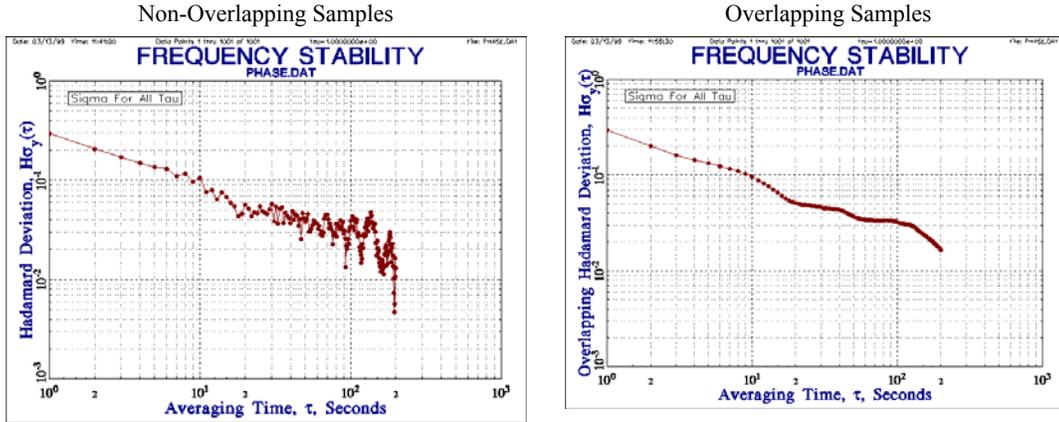
$$H\sigma_y^2(\tau) = \frac{1}{6(M-2)} \sum_{i=1}^{M-2} (y_{i+2} - 2y_{i+1} + y_i)^2 \qquad H\sigma_y^2(\tau) = \frac{1}{6m^2(M-3m+1)} \sum_{j=1}^{M-3m+1} \sum_{i=j}^{j+m-1} (y_{i+2m} - 2y_{i+m} + y_i)^2$$

Figure 8. The reduction in variability by using overlapping samples in calculating the Hadamard deviation.

5.2.4.   Overlapping Allan Variance

The fully overlapping Allan variance, or AVAR, is a form of the normal Allan variance, $\sigma^2_y(\tau)$, that makes maximum use of a data set by forming all possible overlapping samples at each averaging time τ. It can be estimated from a set of M frequency measurements for averaging time $\tau = m\tau_0$, where m is the averaging factor and τ_0 is the basic measurement interval, by the expression

> The overlapped Allan deviation is the most common measure of time-domain frequency stability. The term AVAR has come to be used mainly for this form of the Allan variance, and ADEV for its square root.

$$\sigma_y^2(\tau) = \frac{1}{2m^2(M-2m+1)} \sum_{j=1}^{M-2m+1} \left\{ \sum_{i=j}^{j+m-1} [y_{i+m} - y_i] \right\}^2 . \tag{10}$$

This formula is seldom used for large data sets because of the computationally intensive inner summation. In terms of phase data, the overlapping Allan variance can be estimated from a set of $N = M + 1$ time measurements as

$$\sigma_y^2 = \frac{1}{2(N-2m)\tau^2} \sum_{i=1}^{N-2m} [x_{i+2m} - 2x_{i+m} + x_i]^2 . \tag{11}$$

Fractional frequency data, y_i, can be first integrated to use this faster formula. The result is usually expressed as the square root, $\sigma_y(\tau)$, the Allan deviation, ADEV. The confidence interval of an overlapping Allan deviation estimate is better than that of a normal Allan variance estimation because, even though the additional overlapping differences are not all statistically independent, they nevertheless increase the number of degrees of freedom and thus improve the confidence in the estimation. Analytical methods are available for calculating the number of degrees of freedom for an estimation of overlapping Allan variance, and using that to establish single- or double-sided confidence intervals for the estimate with a certain confidence factor, based on Chi-squared statistics.

Sample variances are distributed according to the expression

$$\chi^2 = \frac{df \cdot s^2}{\sigma^2},$$ (12)

where χ^2 is the Chi-square, s^2 is the sample variance, σ^2 is the true variance, and df is the number of degrees of freedom (not necessarily an integer). For a particular statistic, df is determined by the number of data points and the noise type.

References for Allan Variance

1. D.W. Allan, "The statistics of atomic frequency standards," *Proc. IEEE*, 54(2): 221-230 (Feb. 1966).
2. D.W. Allan, "Allan variance," **http://www.allanstime.com** [2008].
3. "Characterization of frequency stability," Nat. Bur. Stand. (U.S.) Tech Note 394 (Oct. 1970).
4. J.A. Barnes, A.R. Chi, L.S. Cutler, D.J. Healey, D.B. Leeson, T.E. McGunigal, J.A. Mullen, Jr., W.L. Smith, R.L. Sydnor, R.F.C. Vessot, and G.M.R. Winkler, "Characterization of frequency stability," *IEEE Trans. Instrum. Meas.*, 20(2): 105-120 (May 1971)
5. J.A. Barnes, "Variances based on data with dead time between the measurements," Natl. Inst. Stand. Technol. Technical Note 1318 (1990).
6. C.A. Greenhall, "Does Allan variance determine the spectrum?" *Proc. 1997 Intl. Freq. Cont. Symp.*, pp. 358-365 (June 1997)/
7. C.A. Greenhall, "Spectral ambiguity of Allan variance," *IEEE Trans. Instrum. Meas.*, 47(3): 623-627 (June 1998).

5.2.5   Modified Allan Variance

The modified Allan variance, Mod $\sigma^2_y(\tau)$, MVAR, is another common time domain measure of frequency stability [1]. It is estimated from a set of M frequency measurements for averaging time $\tau = m\tau_0$, where m is the averaging factor and τ_0 is the basic measurement interval, by the expression

> Use the modified Allan deviation to distinguish between white and flicker PM noise.

$$Mod\sigma^2_y(\tau) = \frac{1}{2m^4(M-3m+2)} \sum_{j=1}^{M-3m+2} \left\{ \sum_{i=j}^{j+m-1} \left(\sum_{k=i}^{i+m-1} [y_{k+m} - y_k] \right) \right\}^2 .$$ (13)

In terms of phase data, the modified Allan variance is estimated from a set of N = M + 1 time measurements as

$$Mod\sigma^2_y(\tau) = \frac{1}{2m^2\tau^2(N-3m+1)} \sum_{j=1}^{N-3m+1} \left\{ \sum_{i=j}^{j+m-1} [x_{i+2m} - 2x_{i+m} + x_i] \right\}^2 .$$ (14)

The result is usually expressed as the square root, Mod $\sigma_y(\tau)$, the modified Allan deviation. The modified Allan variance is the same as the normal Allan variance for m = 1. It includes an additional phase averaging operation, and has the advantage of being able to distinguish between white and flicker PM noise. The confidence interval of a modified Allan deviation determination is also dependent on the noise type, but is often estimated as $\pm\sigma_y(\tau)/\sqrt{N}$.

References for Modified Allan Variance

1. D.W. Allan and J.A. Barnes, "A modified Allan variance with increased oscillator characterization ability," *Proc. 35th Freq. Cont. Symp.* pp. 470-474 (May 1981).
2. P. Lesage and T. Ayi, "Characterization of frequency stability: Analysis of the modified Allan variance and properties of its estimate," *IEEE Trans. Instrum. Meas.*, 33(4): 332-336 (Dec. 1984).
3. C.A. Greenhall, "Estimating the modified Allan variance," *Proc. IEEE 1995 Freq. Cont. Symp.*, pp. 346-353 (May 1995).
4. C.A. Greenhall, "The third-difference approach to modified Allan variance," *IEEE Trans. Instrum. Meas.*, 46(3): 696-703 (June 1997).

5.2.6. Time Variance

The time Allan variance, TVAR, with square root TDEV, is a measure of time stability based on the modified Allan variance [1]. It is defined as

> Use the time deviation to characterize the time error of a time source (clock) or distribution system.

$$\sigma_x^2(\tau) = \left(\frac{\tau^2}{3}\right) \cdot Mod\,\sigma_y^2(\tau) \,. \tag{15}$$

In simple terms, TDEV is MDEV whose slope on a log-log plot is transposed by +1 and normalized by √3. The time Allan variance is equal to the standard variance of the time deviations for white PM noise. It is particularly useful for measuring the stability of a time distribution network.

It can be convenient to include TDEV information on a MDEV plot by adding lines of constant TDEV, as shown in Figure 9:
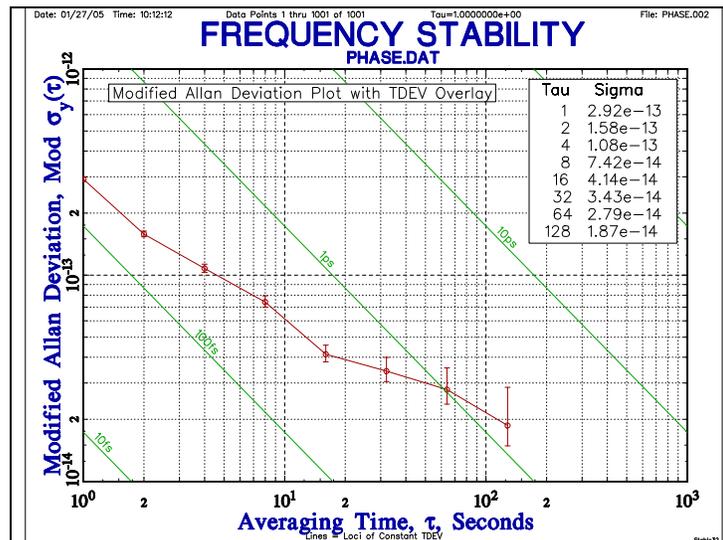
Figure 9. Plot of MDEV with lines of constant TDEV.

References for Time Variance

1.   D.W. Allan, D.D. Davis, J. Levine, M.A. Weiss, N. Hironaka, and D. Okayama, "New inexpensive frequency calibration service from Natl. Inst. Stand. Technol.," *Proc. 44th Freq. Cont. Symp.*, pp. 107-116 (June 1990).
2.   D.W. Allan, M.A. Weiss, and J.L. Jespersen, "A frequency-domain view of time-domain characterization of clocks and time and frequency distribution systems," *Proc. 45th Freq. Cont. Symp.*, pp. 667-678 (May 1991).

5.2.7.   Time Error Prediction

The time error of a clock driven by a frequency source is a relatively simple function of the initial time offset, the frequency offset, and the subsequent frequency drift, plus the effect of noise, as shown in the following expression:

> The time error of a clock can be predicted from its time and frequency offsets, frequency drift, and noise.

$$\Delta T = T_o + (\Delta f/f) \cdot t + \frac{1}{2} D \cdot t^2 + \sigma_x(t), \qquad\qquad (16)$$

where ΔT is the total time error, T_o is the initial synchronization error, $\Delta f/f$ is the sum of the initial and average environmentally induced frequency offsets, D is the frequency drift (aging rate), and $\sigma_x(t)$ is the root-mean-square (rms) noise-induced time deviation. For consistency, units of dimensionless fractional frequency and seconds should be used throughout.

Because of the many factors, conditions, and assumptions involved, and their variability, clock error prediction is seldom easy or exact, and it is usually necessary to generate a timing error budget.

*   Initial Synchronization

The effect of an initial time (synchronization) error, T_o, is a constant time offset due to the time reference, the finite measurement resolution, and measurement noise. The measurement resolution and noise depends on the averaging time.

*   Initial Syntonization

The effect of an initial frequency (syntonization) error, $\Delta f/f$, is a linear time error. Without occasional resyntonization (frequency recalibration), frequency aging can cause this to be the biggest contributor toward clock error for many frequency sources (e.g., quartz crystal oscillators and rubidium gas cell standards). Therefore, it can be important to have a means for periodic clock syntonization (e.g., GPS or cesium beam standard). In that case, the syntonization error is subject to uncertainty due to the frequency reference, the measurement and tuning resolution, and noise considerations. The measurement noise can be estimated by the square root of the sum of the Allan variances of the clock and reference over the measurement interval. The initial syntonization should be performed, to the greatest extent possible, under the same environmental conditions (e.g., temperature) as expected during subsequent operation.

*   Environmental Sensitivity

After initial syntonization, environmental sensitivity is likely to be the largest contributor to time error. Environmental frequency sensitivity obviously depends on the properties of the device and its operating conditions. When performing a frequency stability analysis, it is important to separate the deterministic environmental sensitivities from the stochastic noise. This requires a good understanding of both the device and its environment.

Reference for Time Error Prediction

D.W. Allan and H. Hellwig, "Time Deviation and Time Prediction Error for Clock Specification, Characterization, and Application", *Proceedings of the Position Location and Navigation Symposium (PLANS)*, 29-36, 1978.

5.2.8. Hadamard Variance

The Hadamard [1] variance is based on the Hadamard transform [2], which was adapted by Baugh as the basis of a time-domain measure of frequency stability [3]. As a spectral estimator, the Hadamard transform has higher resolution than the Allan variance, since the equivalent noise bandwidth of the Hadamard and Allan spectral windows are $1.2337N^{-1}\tau^{-1}$ and $0.476\tau^{-1}$, respectively [4]. For the purposes of time-domain frequency stability characterization, the most important advantage of the Hadamard variance is its insensitivity to linear frequency drift, making it particularly useful for the analysis of rubidium atomic clocks [5,6]. It has also been used as one of the components of a time-domain multivariance analysis [7], and is related to the third structure function of phase noise [8].

> Use the Hadamard variance to characterize frequency sources with divergent noise and/or frequency drift.

Because the Hadamard variance examines the second difference of the fractional frequencies (the third difference of the phase variations), it converges for the Flicker Walk FM ($\alpha = -3$) and Random Run FM ($\alpha = -4$) power-law noise types. It is also unaffected by linear frequency drift.

For frequency data, the Hadamard variance is defined as:

$$H\sigma_y^2(\tau) = \frac{1}{6(M-2)} \sum_{i=1}^{M-2} [y_{i+2} - 2y_{i+1} + y_i]^2 , \qquad (17)$$

where y_i is the ith of M fractional frequency values at averaging time τ.

For phase data, the Hadamard variance is defined as:

$$H\sigma_y^2(\tau) = \frac{1}{6\tau^2(N-3m)} \sum_{i=1}^{N-3} [x_{i+3} - 3x_{i+2} + 3x_{i+1} - x_i]^2 , \qquad (18)$$

where x_i is the ith of $N = M + 1$ phase values at averaging time τ.

Like the Allan variance, the Hadamard variance is usually expressed as its square-root, the Hadamard deviation, HDEV or $H\sigma_y(\tau)$.

5.2.9. Overlapping Hadamard Variance

In the same way that the overlapping Allan variance makes maximum use of a data set by forming all possible fully overlapping 2-sample pairs at each averaging time τ the overlapping Hadamard variance uses all 3-sample combinations [9]. It can be estimated from a set of M frequency measurements for averaging time $\tau = m\tau_0$ where m is the averaging factor and τ_0 is the basic measurement interval, by the expression:

> The overlapping Hadamard variance provides better confidence than the non-overlapping version.

$$H\sigma_y^2(\tau) = \frac{1}{6m^2(M-3m+1)\tau^2} \sum_{j=1}^{M-3m+1} \left\{ \sum_{i=j}^{j+m-1} [y_{i+2m} - 2y_{i+m} + y_i] \right\}^2 , \tag{19}$$

where y_i is the ith of M fractional frequency values at each measurement time.

In terms of phase data, the overlapping Hadamard variance can be estimated from a set of N = M + 1 time measurements as:

$$H\sigma_y^2(\tau) = \frac{1}{6(N-3m)\tau^2} \sum_{i=1}^{N-3m} [x_{i+3m} - 3x_{i+2m} + 3x_{i+m} - x_i]^2 , \tag{20}$$

where x_i is the ith of N = M + 1 phase values at each measurement time.

Computation of the overlapping Hadamard variance is more efficient for phase data, where the averaging is accomplished by simply choosing the appropriate interval. For frequency data, an inner averaging loop over m frequency values is necessary. The result is usually expressed as the square root, $H\sigma_y(\tau)$, the Hadamard deviation, HDEV. The expected value of the overlapping statistic is the same as the normal one described above, but the confidence interval of the estimation is better. Even though not all the additional overlapping differences are statistically independent, they nevertheless increase the number of degrees of freedom and thus improve the confidence in the estimation. Analytical methods are available for calculating the number of degrees of freedom for an overlapping Allan variance estimation, and that same theory can be used to establish reasonable single- or double-sided confidence intervals for an overlapping Hadamard variance estimate with a certain confidence factor, based on Chi-squared statistics.

Sample variances are distributed according to the expression:

$$\chi^2(p, df) = \frac{df \cdot s^2}{\sigma^2} , \tag{21}$$

where χ^2 is the Chi-square value for probability p and degrees of freedom df, s^2 is the sample variance, σ^2 is the true variance, and df is the number of degrees of freedom (not necessarily an integer). The df is determined by the number of data points and the noise type. Given the df, the confidence limits around the measured sample variance are given by:

$$\sigma_{min}^2 = \frac{(s^2 \cdot df)}{\chi^2(p, df)} \text{ and } \sigma_{max}^2 = \frac{(s^2 \cdot df)}{\chi^2(1-p, df)} . \tag{22}$$

5.2.10.   Modified Hadamard Variance

By similarity to the modified Allan variance, a modified version of the Hadamard variance can be defined [15] that employs averaging of the phase data over the m adjacent samples that define the analysis $\tau = m \cdot \tau_0$. In terms of phase data, the three-sample modified Hadamard variance is defined as:

$$Mod\sigma_H^2(\tau) = \frac{\sum_{j=1}^{N-4m+1} \left\{ \sum_{i=j}^{j+m-1} [x_i - 3x_{i+m} + 3x_{i+2m} - x_{i+3m}] \right\}^2}{6m^2\tau^2[N-4m+1]} , \tag{23}$$

where N is the number of phase data points x_i at the sampling interval τ_0, and m is the averaging factor, which can extend from 1 to $\lfloor N/4 \rfloor$. This is an unbiased estimator of the modified Hadamard variance, MHVAR. Expressions for the equivalent number of χ^2 degrees of freedom (edf) required to set MHVAR confidence limits are available in [2].

Clock noise (and other noise processes) can be described in terms of power spectral density, which can be modeled as a power law function $S \propto f^{\alpha}$, where f is Fourier frequency and α is the power law exponent. When a variance such as MHVAR is plotted on log-log axes versus averaging time, the various power law noises correspond to particular slopes μ. MHVAR was developed in Reference [15] for determining the power law noise type of Internet traffic statistics, where it was found to be slightly better for that purpose than the modified Allan variance, MVAR, when there were a sufficient number of data points. MHVAR could also be useful for frequency stability analysis, perhaps in cases where it was necessary to distinguish between short-term white and flicker PM noise in the presence of more divergent ($\alpha = -3$ and -4) flicker walk and random run FM noises. The Mod $\sigma^2_H(\tau)$ log-log slope μ is related to the power law noise exponent by $\mu = -3 - \alpha$.

The modified Hadamard variance concept can be generalized to subsume AVAR, HVAR, MVAR, MHVAR, and MHVARs using higher-order differences:

$$Mod\sigma^2{}_{H,d}(\tau) = \frac{\sum\limits_{j=1}^{N-(d+1)m+1}\left\{\sum\limits_{i=j}^{j+m-1}\sum\limits_{k=0}^{d}\binom{d}{k}(-1)^k x_{i+km}\right\}^2}{d!\,m^2\tau^2\left[N-(d+1)m+1\right]}, \tag{24}$$

where d = phase differencing order; $d = 2$ corresponds to MAVAR, $d = 3$ to MHVAR; higher-order differencing is not commonly used in the field of frequency stability analysis. The unmodified, nonoverlapped AVAR and HVAR variances are given by setting $m = 1$. The allowable power law exponent for convergence of the variance is equal to $\alpha > 1 - 2d$, so the second difference Allan variances can be used for $\alpha > -3$ and the third difference Hadamard variances for $\alpha > -5$.

Confidence intervals for the modified Hadamard variance can be determined by use of the edf values of reference [16].

References for Hadamard Variance

1. Jacques Saloman Hadamard (1865−1963), French mathematician.
2. W.K. Pratt, J. Kane, and H.C. Andrews, "Hadamard transform image coding," *Proc. IEEE*, 57(1): 38-67 (Jan. 1969).
3. R.A. Baugh, "Frequency modulation analysis with the Hadamard variance," *Proc. Freq. Cont. Symp*., pp. 222-225 (June 1971).
4. K. Wan, E. Visr, and J. Roberts, "Extended variances and autoregressive moving average algorithm for the measurement and synthesis of oscillator phase noise," *Proc. 43rd Freq. Cont. Symp*., pp.331-335 (June 1989).
5. S.T. Hutsell, "Relating the Hamamard variance to MCS Kalman filter clock estimation," *Proc. 27th PTTI Mtg*., pp. 291-302 (Dec. 1995).
6. S.T. Hutsell, "Operational use of the Hamamard variance in GPS," *Proc. 28th PTTI Mtg*., pp. 201-213 (Dec. 1996).
7. T. Walter, "A multi-variance analysis in the time domain," *Proc. 24th PTTI Mtg*., pp. 413-424 (Dec. 1992).
8. J. Rutman, "Oscillator specifications: A review of classical and new ideas," *1977 IEEE Intl. Freq. Cont. Symp*., pp. 291-301 (June 1977).
9. This expression for the overlapping Hadamard variance was developed by the author at the suggestion of G. Dieter and S.T. Hutsell.
10. Private communication, C. Greenhall to W. Riley, 5/7/99.
11. B. Picinbono, « Processus à accroissements stationnaires » *Ann. des telecom*, 30(7-8): 211-212 (July-Aug. 1975).
12. E. Boileau and B. Picinbono, "Statistical study of phase fluctuations and oscillator stability," *IEEE Trans. Instrum. Meas*., 25(1): 66-75 (March 1976).
13. D.N. Matsakis and F.J. Josties, "Pulsar-appropriate clock statistics," *Proc. 28th PTTI Mtg*., pp. 225-236 (Dec. 1996).
14. Chronos Group, Frequency Measurements and Control, Sect. 3.3.3, Chapman & Hall, London, ISBN 0-412-48270-3 (1994).
15. S. Bregni and L. Jmoda, "Improved estimation of the Hurst parameter of long-range dependent traffic using the modified Hadamard variance," *Proc. IEEE ICC* (June 2006).
16. C.A. Greenhall and W.J. Riley, "Uncertainty of stability variances based on finite differences," *Proc. 35th PTTI Mtg*. (Dec. 2003).

5.2.11.　　Total Variance

The total variance, TOTVAR, is similar to the two-sample or Allan variance and has the same expected value, but offers improved confidence at long averaging times [1-5]. The work on total variance began with the realization that the Allan variance can "collapse" at long averaging factors because of symmetry in the data. An early idea was to shift the data by 1/4 of the record length and average the two resulting Allan variances. The next step was to wrap the data in a circular fashion and calculate the

> The total variance offers improved confidence at large averaging factor by extending the data set by reflection at both ends.

average of all the Allan variances at every basic measurement interval, τ_o. This technique is very effective in improving the confidence at long averaging factors but requires end matching of the data. A further improvement of the total variance concept was to extend the data by reflection, first at one end of the record and then at both ends. This latest technique, called TOTVAR, gives a very significant confidence advantage at long averaging times, exactly decomposes the classical standard variance [6], and is an important new general statistical tool. TOTVAR is defined for phase data as:

$$Tot \, \text{var}(\tau) = \frac{1}{2\tau^2(N-2)} \sum_{i=2}^{N-1} \left[x^*_{i-m} - 2x^*_i + x^*_{i+m} \right]^2 , \qquad (25)$$

where $\tau = m\tau_o$, and the N phase values x measured at $\tau = \tau_o$ are extended by reflection about both endpoints to form a virtual sequence x^* from $i = 3–N$ to $i = 2N–2$ of length $3N–4$. The original data are in the center of x^* with $i = 1$ to N and $x^* = x$. The reflected portions added at each end extend from $j = 1$ to $N–2$ where $x^*_{1-j} = 2x_1 – x_{1+j}$ and $x^*_{N+j} = 2x_N – x_{N-j}$.

Totvar can also be defined for frequency data as:

$$Tot\,\mathrm{var}(\tau) = \frac{1}{2(M-1)} \sum_{i=1}^{M-1} \left[y^*_{i+j+1} - y^*_{i+j} \right]^2 , \tag{26}$$

where the $M = N–1$ fractional frequency values, y, measured at $\tau = \tau_o$ (N phase values) are extended by reflection at both ends to form a virtual array y^*. The original data are in the center, where $y^*_I = y_i$ for $i = 1$ to M, and the extended data for $j = 1$ to $M–1$ are equal to $y^*_{1-j} = y_j$ and $y^*_{M+1} = y_{M+1-j}$.

The result is usually expressed as the square root, $\sigma_{total}(\tau)$, the total deviation, TOTDEV. When calculated by use of the doubly reflected method described above, the expected value of TOTVAR is the same as AVAR for white and flicker PM or white FM noise. Bias corrections of the form $1/[1–a(\tau/T)]$, where T is the record length, need to be applied for flicker and random walk FM noise, where $a = 0.481$ and 0.750, respectively.

The number of equivalent χ^2 degrees of freedom for TOTVAR can be estimated for white FM, flicker FM and random walk FM noise by the expression $b(T/\tau)–c$, where $b = 1.500$, 1.168 and 0.927, and $c = 0$, 0.222 and 0.358, respectively. For white and flicker PM noise, the edf for a total deviation estimate is the same as that for the overlapping ADEV with the number of χ^2 degrees of freedom increased by 2.

References for Total Variance

1. D.A. Howe, "An extension of the Allan variance with increased confidence at long term," *Proc. 1995 IEEE Int. Freq. Cont. Symp.*, pp. 321-329 (June 1995).
2. D.A. Howe and K.J. Lainson, "Simulation study using a new type of sample variance," *Proc. 1995 PTTI Mtg.*, pp. 279-290 (Dec. 1995).
3. D.A. Howe and K.J. Lainson, "Effect of drift on TOTALDEV," *Proc. 1996 Intl. Freq. Cont. Symp.*, pp. 883-889 (June 1996).
4. D.A. Howe, "Methods of improving the estimation of long-term frequency variance," *Proc. 1997 European Frequency and Time Forum*, pp. 91-99 (March 1997).
5. D.A. Howe and C.A. Greenhall, "Total variance: A progress report on a new frequency stability characterization," *Proc. 1997 PTTI Mtg.*, pp. 39-48 (Dec. 1997).
6. D.B. Percival and D.A. Howe, "Total variance as an exact analysis of the sample variance," *Proc. 1997 PTTI Mtg.*, pp. 97-105 (Dec. 1997).
7. C.A. Greenhall, D.A. Howe, and D.B. Percival, "Total variance, an estimator of long-term frequency stability," *IEEE Trans. Ultrason. Ferroelect. Freq. Cont.*, 46(5): 1183-1191 (Sept. 1999).
8. D. Howe and T. Peppler, "Definitions of total estimators of common time-domain variances," *Proc. 2001 IEEE Intl. Freq. Cont. Symp.* 127-132 (June 2001).

5.2.12.　　**Modified Total Variance**

The modified total variance, MTOT, is another new statistic for the analysis of frequency stability. It is similar to the modified Allan variance, MVAR, and has the same expected value, but offers improved confidence at long averaging times. It uses the same phase averaging technique as MVAR to distinguish between white and flicker PM noise processes.

> The modified total variance combines the features of the modified Allan and total variances.

A calculation of MTOT begins with an array of N phase data points (time deviates, x_i) with sampling period τ_o that are to be analyzed at averaging time $\tau = m\tau_0$. MTOT is computed from a set of N − 3m + 1 subsequences of 3m points. First, a linear trend (frequency offset) is removed from the subsequence by averaging the first and last halves of the subsequence and dividing by half the interval. Then the offset-removed subsequence is extended at both ends by uninverted, even reflection. Next the modified Allan variance is computed for these 9m points. Finally, these steps are repeated for each of the N − 3m + 1 subsequences, calculating MTOT as their overall average. These steps, similar to those for MTOT, but acting on fractional frequency data, are shown in Figure 10.
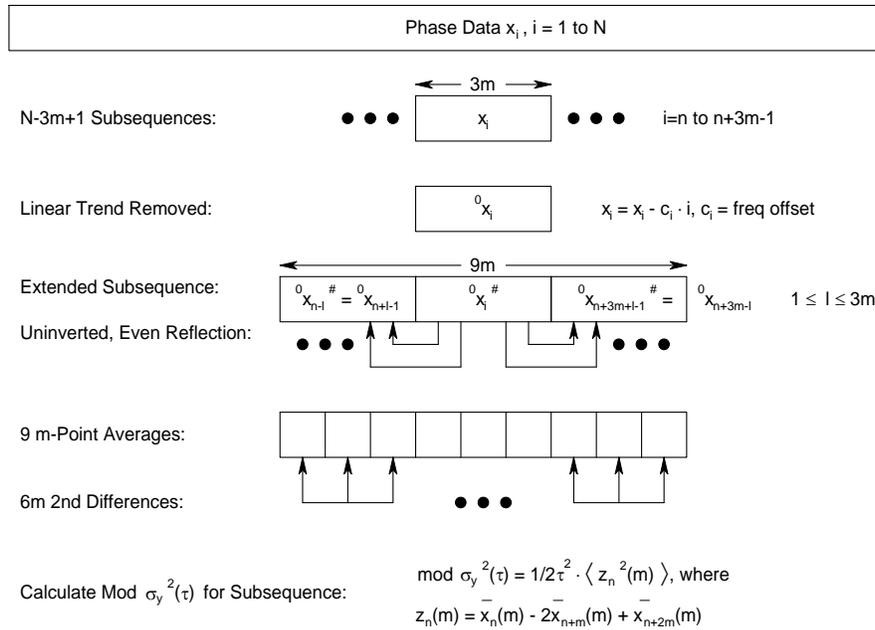
Figure 10.　Steps similar to calculation of MTOT on fractional frequency data.

Computationally, the MTOT process requires three nested loops:

- An outer summation over the N − 3m + 1 subsequences. The 3m-point subsequence is formed, its linear trend is removed, and it is extended at both ends by uninverted, even reflection to 9m points.
- An inner summation over the 6m unique groups of m-point averages from which all possible fully overlapping second differences are used to calculate MVAR.
- A loop within the inner summation to sum the phase averages for three sets of m points.

The final step is to scale the result according to the sampling period, τ_0, averaging factor, m, and number of points, N. Overall, this can be expressed as:

$$Mod\ Tot\ \mathrm{var}(\tau) = \frac{1}{2(m\tau_0)^2(N-3m+1)} \sum_{n=1}^{N-3m-1} \left\{ \frac{1}{6m} \sum_{i=n-3m}^{n+3m-1} \left[{}^0z_i^{\#}(m) \right]^2 \right\}, \tag{27}$$

25

where the $^0z_i^{\#}(m)$ terms are the phase averages from the triply extended subsequence, and the prefix ^{0} denotes that the linear trend has been removed. At the largest possible averaging factor, $m = N/3$, the outer summation consists of only one term, but the inner summation has 6m terms, thus providing a sizable number of estimates for the variance.

Reference for Modified Total Variance

D.A. Howe and F. Vernotte, "Generalization of the Total Variance Approach to the Modified Allan Variance," *Proc. 31ˢᵗ PTTI Meeting*, pp. 267-276, Dec. 1999.

5.2.13.   Time Total Variance

The time total variance, TTOT, is a similar measure of time stability, based on the modified total variance. It is defined as

> The time total variance is a measure of time stability based on the modified total variance.

$$\sigma_x^2(\tau) = \frac{\tau^3}{3} \cdot \mathrm{Mod}\; \sigma_{total(\tau)}^2 \qquad (28)$$

5.2.14.   Hadamard Total Variance

The Hadamard total variance, HTOT, is a total version of the Hadamard variance. As such, it rejects linear frequency drift while offering improved confidence at large averaging factors.

> The Hadamard total variance combines the features of the Hadamard and total variances by rejecting linear frequency drift, handling more divergent noise types, and providing better confidence at large averaging factors.

An HTOT calculation begins with an array of N fractional frequency data points, y_i with sampling period τ_0 that are to be analyzed at averaging time $\tau = m\,\tau_0$. HTOT is computed from a set of $N - 3m + 1$ subsequences of 3m points. First, a linear trend (frequency drift) is removed from the subsequence by averaging the first and last halves of the subsequence and dividing by half the interval. Then the drift-removed subsequence is extended at both ends by uninverted, even reflection. Next the Hadamard variance is computed for these 9m points. Finally, these steps are repeated for each of the $N - 3m + 1$ subsequences, calculating HTOT as their overall average. These steps are shown in Figure11.

Fractional Frequency Data y, i = 1 to N

N-3m+1 Subsequences: ... y_i ... i=n to n+3m-1

3m

Linear Freq Drift Removed: ^0y_i $y_i = y_i - c_i \cdot i$, c_i = freq drift

9m

Extended Subsequence: $^0y_{n-l}$ # = $^0y_{n+l-1}$ ^0y_i # $^0y_{n+3m+l-1}$ # = $^0y_{n+3m-l}$ $1 \le l \le 3m$

Uninverted, Even Reflection: ...   ...

9 m-Point Averages:

6m 2nd Differences: ...

Calculate $Had\sigma_y^2(\tau)$ for Subsequence: $Had\sigma_y^2(\tau) = 1/6 \cdot \langle z_n^2(m) \rangle$, where
$$z_n(m) = \bar{y}_n(m) - 2\bar{y}_{n+m}(m) + \bar{y}_{n+2m}(m)$$

Then Find HTOT as Average of Subestimates

Figure 11. Steps to calculate Hadamard Total Variance.

Computationally, the HTOT process requires three nested loops:

- An outer summation over the $N - 3m + 1$ subsequences. The 3m-point subsequence is formed, its linear trend is removed, and it is extended at both ends by uninverted, even reflection to 9m points.
- An inner summation over the 6m unique groups of m-point averages from which all possible fully overlapping second differences are used to calculate HVAR.
- A loop within the inner summation to sum the frequency averages for three sets of m points.

The final step is to scale the result according to the sampling period, τ_0, averaging factor, m, and number of points, N. Overall, this can be expressed as:

$$TotalH\sigma_y^2\left(m,\tau_0,N\right) = \frac{1}{6(N-3m+1)} \sum_{n=1}^{N-3m+1}\left(\frac{1}{6m}\sum_{i=n-3m}^{n+3m-1}(H_i(m))^2\right), \tag{29}$$

where the $H_i(m)$ terms are the $z_n(m)$ Hadamard second differences from the triply extended, drift-removed subsequences. At the largest possible averaging factor, m = N/3, the outer summation consists of only one term, but the inner summation has 6m terms, thus providing a sizable number of estimates for the variance. The Hadamard total variance is a biased estimator of the Hadamard variance, so a bias correction is required that is dependent on the power law noise type and number of samples.

The following plots shown the improvement in the consistency of the overlapping Hadamard deviation results compared with the normal Hadamard deviation, and the extended averaging factor range provided by the Hadamard total deviation [10].
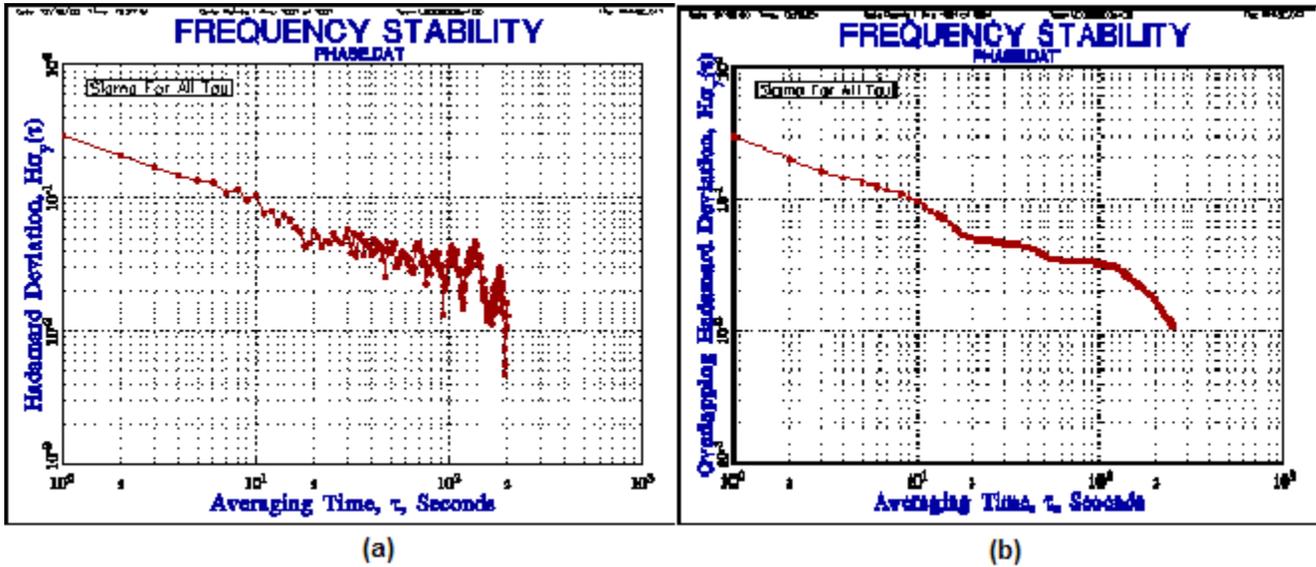
Figure 12. (a) Hadamard Deviation, (b) Overlapping Hadamard Deviation.
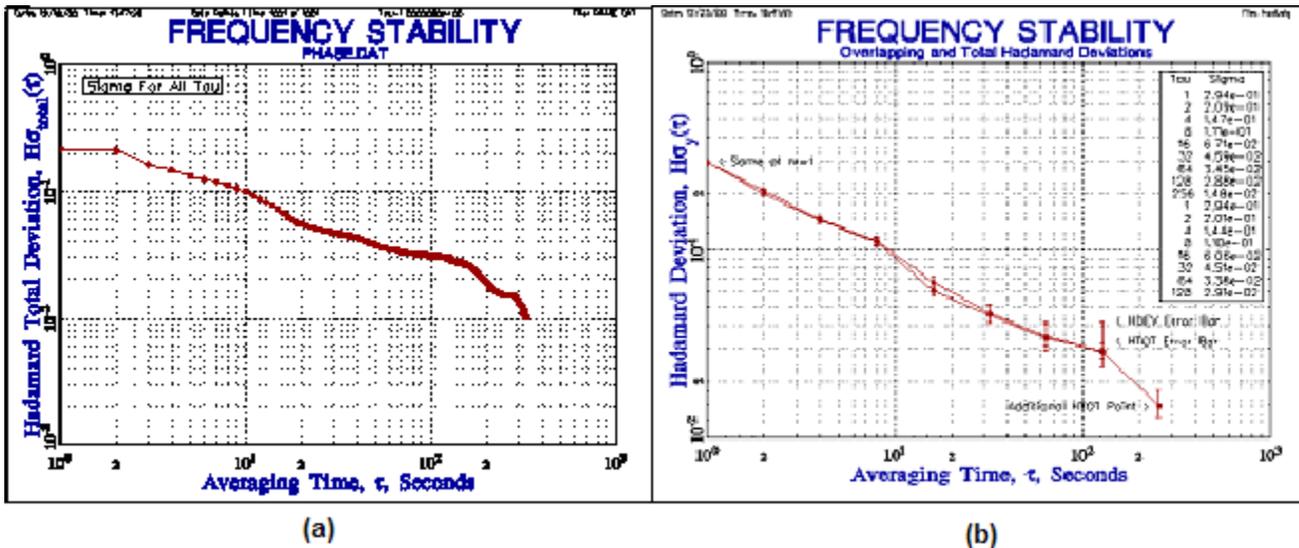
Figure 13. (a) Hadamard Total Deviation, (b) Overlapping & Total Hadamard Deviations

A comparison of the overlapping and total Hadamard deviations shows the tighter error bars of the latter, allowing an additional point to be shown at the longest averaging factor.

The Hadamard variance may also be used to perform a frequency domain (spectral) analysis because it has a transfer function that is a close approximation to a narrow rectangle of spectral width $1/(2 \cdot N \cdot \tau_0)$, where N is the number of samples, and τ_0 is the measurement time [3]. This leads to a simple expression for the spectral density of the fractional frequency fluctuations $S_y(f) \approx 0.73 \cdot \tau_0 \cdot H\sigma^2_y(\tau) / N$, where $f = 1/ (2 \cdot \tau_0)$, which can be particularly useful at low Fourier frequencies.

The Picinbono variance is a similar three-sample statistic. It is identical to the Hadamard variance except for a factor of 2/3 [4]. Sigma-z is another statistic that is similar to the Hadamard variance that has been applied to the study of pulsars [5].

28

It is necessary to identify the dominant power law noise type as the first step in determining the estimated number of chi-squared degrees of freedom for the Hadamard statistics so their confidence limits can be properly set [6]. Because the Hadamard variances can handle the divergent flicker walk FM and random run FM power law noises, techniques for those noise types must be included. Noise identification is particularly important for applying the bias correction to the Hadamard total variance.

References for Hadamard Total Variance

1. D.A. Howe, R. Beard, C.A. Greenhall, F. Vernotte, and W.J. Riley, "A total estimator of the Hadamard function used for GPS operations," *Proc. 32nd PTTI Mtg.*, pp. 255-267, Nov. 2000.
2. D.A. Howe, R. Beard, C.A. Greenhall, F. Vernotte, and W.J. Riley, "Total Hadamard variance: Application to clock steering by Kalman filtering," *Proc. 2001 European Freq. and Time Forum*, pp. 423-427 (Mar. 2001).
3. Chronos Group, Frequency measurement and control, Section 3.3.3, Chapman & Hall, London, ISBN 0-412-48270-3 (1994).
4. B. Picinbono, « Processus à accroissements stationnaires » *Ann. des telecom*, 30(7-8): 211-212 (July-Aug. 1975).
5. D.N. Matsakis and F.J. Josties, "Pulsar-appropriate clock statistics," *Proc. 28th PTTI Mtg.*, pp. 225-236 (Dec. 1996).
6. D.A. Howe, R.L. Beard, C.A. Greenhall, F. Vernotte, W.J. Riley, and T.K. Peppler, "Enhancements to GPS operations and clock evaluations using a total Hadamard deviation," *IEEE Trans. Ultrason. Ferroelect. Freq. Cont.*, 52: 1253-1261 (Aug. 2005).

5.2.15. Thêo1

Thêo1 is a new class of statistics which mimic the properties of the Allan variance (AVAR) while covering a larger range of averaging times, 10 to N−2 for Thêo1 *vs.* 1 to (N−1)/2 for AVAR [1]. It provides improved confidence and the ability to obtain a result for a maximum averaging time equal to 75 % of the record length.

> Thêo1 is a two-sample variance with improved confidence and extended averaging factor range.

Thêo1 [1] is defined as follows:

$$Theo1(m, \tau_0, N) = \frac{1}{(N-m)(m\tau_0)^2} \sum_{i=1}^{N-m} \sum_{\delta=0}^{m/2-1} \frac{1}{m/2-\delta} \left[(x_i - x_{i-\delta+m/2}) + (x_{i+m} - x_{i+\delta+m/2}) \right]^2, \qquad (30)$$

where m = averaging factor, τ_0 = measurement interval, and N = number of phase data points, for m even, and $10 \leq m \leq N - 1$. It consists of N - m outer sums over the number of phase data points −1, and $m/2$ inner sums. Thêo1 is the rms of frequency differences averaged over an averaging time $\tau = 0.75 (m - 1)\tau_0$.

A schematic for a Thêo1 calculation is shown in Figure 14. This example is for eleven phase samples ($N = 11$) at the largest possible averaging factor ($m = 10$).
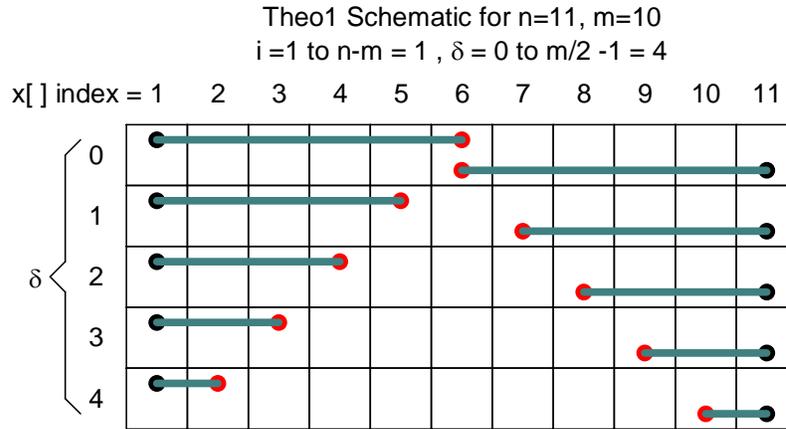
Figure 14. A schematic for Thêo1 calculation.

The single outer summation ($i = 1$ to 1) at the largest possible averaging factor consists of $m/2 = 5$ terms, each with two phase differences. These terms are scaled by their spans $m/2 - \delta = 5$ thru 1 so that they all have equal weighting. A total of 10 terms contribute to the Theo1 statistic at this largest-possible averaging factor. The averaging time, τ, associated with a Thêo1 value is $\tau = 0.75 \cdot m \cdot \tau_0$, where τ_0 is the measurement interval. Thêo1 has the same expected value as the Allan variance for white FM noise, but provides many more samples that provide improved confidence and the ability to obtain a result for a maximum τ equal to three-fourths of the record length, T. Thêo1 is a biased estimator of the Allan variance, AVAR, for all noise types except white FM noise, and it therefore requires the application of a bias correction. Reference [2] contains the preferred expression for determining the Thêo1 bias as a function of noise type and averaging factor:

$$\text{Thêo1 Bias} = \frac{\text{Avar}}{\text{Thêo1}} = a + \frac{b}{m^c} ,$$
(31)

where m is the averaging factor and the constants a, b and c are given in Table 2. Note that the effective tau for a Thêo1 estimation is $\tau = 0.75 \cdot m \cdot \tau_0$, where τ_0 is the measurement interval.

Table 2. Thêo1 bias parameters.

Noise	Alpha	a	b	c
RW FM	−2	2.70	−1.53	0.85
F FM	−1	1.87	−1.05	0.79
W FM	0	1.00	0.00	0.00
F PM	1	0.14	0.82	0.30
W PM	2	0.09	0.74	0.40

Empirical formulae have been developed [1] for the number of equivalent χ^2 degrees of freedom for the Thêo1 statistic, as shown in Table 3:

Table 3. Thêo1 EDF Formulae

Noise	EDF
RW FM	$\left(\dfrac{4.4N-2}{2.9r}\right)\left(\dfrac{\left(4.4N-1\right)^2-8.6r(4.4N-1)+11.4r^2)}{(4.4N-3)^2}\right)$
F FM	$\left(\dfrac{2N^2-1.3Nr-3.5r}{Nr}\right)\left(\dfrac{r^3}{r^3+2.3}\right)$
W FM	$\left[\dfrac{4.1N+0.8}{r}-\dfrac{3.1N+6.5}{N}\right]\left(\dfrac{r^{3/2}}{r^{3/2}+5.2}\right)$
F PM	$\left(\dfrac{4.798N^2-6.374Nr+12.387r}{(r+36.6)^{1/2}(N-r)}\right)\left(\dfrac{r}{r+0.3}\right)$
W PM	$\left(\dfrac{0.86(N+1)(N-4r/3)}{N-r}\right)\left(\dfrac{r}{r+1.14}\right)$
where $r = 0.75m$, and with the condition $\tau_0 \le T/10$.	

5.2.16. ThêoH

Thêo1 has the same expected value as the Allan variance if bias is removed [2]. It is useful to combine a bias-removed version of Thêo1, called ThêoBR, with AVAR to produce a composite stability plot. The composite is called "ThêoH" which is short for "hybrid-ThêoBR" [3]. ThêoH is the best statistic available for estimating the stability level and type of noise of a frequency source, particularly at large averaging times and with a mixture of noise types [4].

> NewThêo1, ThêoBR, and ThêoH are versions of Thêo1 that provide bias removal and combination with the Allan variance.

The NewThêo1 algorithm of Reference [2] provides a method of automatic bias correction for a Thêo1 estimation based on the average ratio of the Allan and Thêo1 variances over a range of averaging factors:

$$\text{NewThêo1}(m,\tau_0,N)=\left[\frac{1}{n+1}\sum_{i=0}^{n}\frac{\text{Avar}(m=9+3i,\tau_0,N)}{\text{Thêo1}(m=12+4i,\tau_0,N)}\right]\text{Thêo1}(m,\tau_0,N), \tag{32}$$

where $n = \left\lfloor \dfrac{N}{30}-3\right\rfloor$, and $\lfloor\ \rfloor$ denotes the floor function.

NewThêo1 was used in Reference [2] to form a composite AVAR/ NewThêo1 result called LONG, which has been superseded by ThêoH (see below).

ThêoBR [3] is an improved bias-removed version of Thêo1 given by

$$\text{ThêoBR}(m,\tau_0,N)=\left[\frac{1}{n+1}\sum_{i=0}^{n}\frac{\text{Avar}(m=9+3i,\tau_0,N)}{\text{Thêo1}(m=12+4i,\tau_0,N)}\right]\text{Thêo1}(m,\tau_0,N), \tag{33}$$

where $n = \left\lfloor \dfrac{N}{6}-3\right\rfloor$, and $\lfloor\ \rfloor$ denotes the floor function.

ThêoBR can determine an unbiased estimate of the Allan variance over the widest possible range of averaging times without explicit knowledge of the noise type. ThêoBR in the equation above is computationally intensive for large data sets, but computation time is significantly reduced by phase averaging with negligible effect on bias removal [5].

ThêoH is a hybrid statistic that combines ThêoBR and AVAR on one plot:

$$\text{ThêoH}(m, \tau_0, N) = \left\{ \begin{array}{l} \text{Avar}\,(m, \tau_0, N) \text{ for } 1 \le m \le \dfrac{k}{\tau_0} \\ \text{ThêoBR}\,(m, \tau_0, N) \text{ for } \dfrac{k}{0.75\,\tau_0} \le m \le N-1,\, m \text{ even} \end{array} \right\}, \qquad (34)$$

where k is the largest available $\tau \le 20\%$ T.

An example of a ThêoH plot is shown in Figure 15:
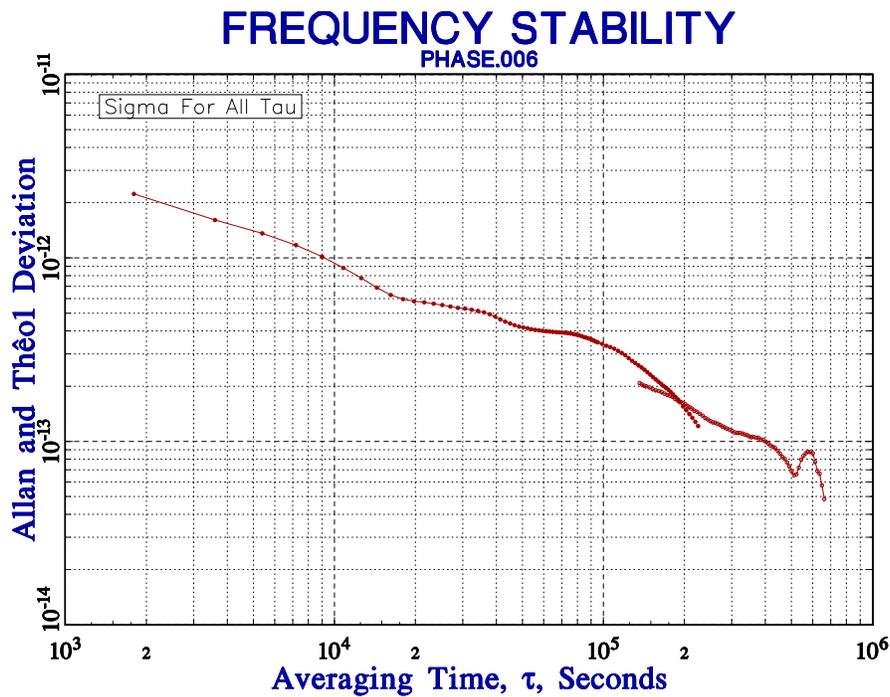
Figure 15. An example plot of ThêoH.

ThêoH is a composite of AVAR and bias-corrected ThêoBR analysis points at a number of averaging times sufficiently large to form a quasi-continuous curve. The data are a set of 1001 simulated phase values measured at 15-minute intervals taken over a period of about 10 days. The AVAR results are able to characterize the stability to an averaging time of about two days, while Thêo1 is able to extend the analysis out to nearly a week, thus providing significantly more information from the same data set. An example of analysis using ThêoH with data from a Cs standard is shown in Section 11.

References for Thêo1, NewThêo1, ThêoBR and ThêoH

1. D.A. Howe and T.K. Peppler, "Very Long-Term Frequency Stability: Estimation Using a Special-Purpose Statistic," *Proceedings of the 2003 IEEE International Frequency Control Symposium*, pp. 233-238, May 2003.
2. D.A. Howe and T.N. Tasset, "Thêo1: Characterization of Very Long-Term Frequency Stability," *Proc. 2004 EFTF.*
3. D.A. Howe, "ThêoH: A Hybrid, High-Confidence Statistic that Improves on the Allan Deviation," *Metrologia* **43** (2006), S322-S331.
4. J. McGee and D.A. Howe, "ThêoH and Allan Deviation as Power-Law Noise Estimators," *IEEE Trans. Ultrasonics, Ferroelectrics and Freq. Contrl.*, Feb. 2007.
5. J. McGee and D.A. Howe, "Fast TheoBR: A method for long data set stability analysis," *IEEE Trans. Ultrasonics, Ferroelectrics and Freq. Contrl.*, to be published, 2008.

5.2.17   MTIE

The maximum time interval error, MTIE, is a measure of the maximum time error of a clock over a particular time interval. This statistic is commonly used in the telecommunications industry. It is calculated by moving an n-point (n = τ/τ_o) window through the phase (time error) data and finding the difference between the maximum and minimum values (range) at each window position. MTIE is the overall maximum of this time interval error over the entire data set:

> MTIE is a measure of clock error commonly used in the tele-communications industry.

$$MTIE(\tau) = Max_{1 \le k \le N-n}\{Max_{k \le i \le k+n}(x_i) - Min_{k \le i \le k+n}(x_i)\}, \qquad (35)$$

where n = 1,2,..., N−1 and N = number of phase data points.

MTIE is a measure of the peak time deviation of a clock and is therefore very sensitive to a single extreme value, transient or outlier. The time required for an MTIE calculation increases geometrically with the averaging factor, n, and can become very long for large data sets (although faster algorithms are available – see [1] below).

The relationship between MTIE and Allan variance statistics is not completely defined, but has been the subject of recent theoretical work [2,3]. Because of the peak nature of the MTIE statistic, it is necessary to express it in terms of a probability level, β, that a certain value is not exceeded.

For the case of white FM noise (important for passive atomic clocks such as the most common rubidium and cesium frequency standards), MTIE can be approximated by the relationship

$$MTIE(\tau, \beta) = k_\beta \cdot \sqrt{h_0 \cdot \tau} = k_\beta \cdot \sqrt{2} \cdot \sigma_y(\tau) \cdot \tau, \qquad (36)$$

where k_β is a constant determined by the probability level, β, as given in Table 4, and h_o is the white FM power-law noise coefficient.

Table 4. Constants β and k_β.

β, %	k_β
95	1.77
90	1.59
80	1.39

The maximum time interval error (MTIE) and rms time interval error (TIE rms) are clock stability measures commonly used in the telecom industry [4, 5]. MTIE is determined by the extreme time deviations within a sliding

window of span τ, and is not as easily related to such clock noise processes as TDEV [2]. MTIE is computationally intensive for large data sets [6].

References for MTIE

1. S. Bregni, "Measurement of maximum time interval error for telecommunications clock stability characterization," *IEEE Trans. Instrum. Meas*., 45(5): 900-906 (Oct. 1996).
2. P. Travella and D. Meo, "The range covered by a clock error in the case of white FM," *Proc. 30th PTTI Mtg*., pp. 49-60 (Dec. 1998).
3. P. Travella, A. Dodone, and S. Leschiutta, "The range covered by a random process and the new definition of MTIE," *Proc. 28th PTTI Mtg*., pp. 119-123 (Dec. 1996).
4. S. Bregni, "Clock Stability Characterization and Measurement in Telecommunications," *IEEE Trans. Instrum. Meas*., 46(6): 1284-1294 (Dec. 1997).
5. G. Zampetti, "Synopsis of timing measurement techniques used in telecommunications," *Proc. 24th PTTI Mtg*., pp. 313-326 (Dec. 1992).
6. S. Bregni and S. Maccabruni, "Fast computation of maximum time interval error by binary decomposition," *IEEE Trans. Instrum. Meas*., 49(6): 1240-1244 (Dec. 2000).
7. M.J. Ivens, "Simulating the Wander accumulation in a SDH synchronisation network," Master's Thesis, University College, London, U.K. (Nov. 1997).

5.2.18.   TIE rms

The rms time interval error, TIE rms, is another clock statistic commonly used by the telecommunications industry. TIE rms is defined by the expression

$$TIE_{rms} = \sqrt{\frac{1}{N-n} \sum_{i=1}^{N-n} \left(x_{i+n} - x_i \right)^2} \ , \tag{37}$$

where n = 1,2,..., N−1 and N = # phase data points.

For no frequency offset, TIE rms is approximately equal to the standard deviation of the fractional frequency fluctuations multiplied by the averaging time. It is therefore similar in behavior to TDEV, although the latter properly identifies divergent noise types.

Reference for TIE rms

S. Bregni, "Clock Stability Characterization and Measurement in Telecommunications," *IEEE Trans. Instrum. Meas*., Vol. 46, No. 6, pp. 1284-1294, Dec. 1997.

5.2.19.   Integrated Phase Jitter and Residual FM

Integrated phase jitter and residual FM are other ways of expressing the net phase or frequency jitter by integrating it over a certain bandwidth. These can be calculated from the amplitudes of the various power law terms.

The power law model for phase noise spectral density (see section 6.1) can be written as

$$S_\phi(f) = K \cdot f^x, \tag{38}$$

where S_ϕ is the spectral density of the phase fluctuations in rad^2/Hz, f is the modulation frequency, K is amplitude in rad^2, and x is the power law exponent. It can be represented as a straight line segment on a plot of $S_\phi(f)$ in dB relative to 1 rad^2/Hz versus log f in hertz. Given two points on the plot (f_1, dB_1) and $f_2, dB_2)$, the values of x and K may be determined by

$$x = \frac{dB_1 - dB_2}{10 \cdot (\log f_1 - \log f_2)} \quad , \tag{39}$$

and

$$K = 10^{\left(\frac{dB_1}{10} - x \log f_1\right)} . \tag{40}$$

The integrated phase jitter can then be found over this frequency interval by

$$\Delta\phi^2 = \int_{f_1}^{f_2} S_\phi(f) \cdot df = \int_{f_1}^{f_2} K \cdot f \cdot df$$

$$\Delta\phi^2 = \frac{K}{x+1}(f_2^{x+1} - f_1^{x+1}) \text{ for } x \neq -1 \tag{41}$$

$$\Delta\phi^2 = K \cdot (\log f_2 - \log f_1) \text{ for } x = 1.$$

It is usually expressed as $\Delta\phi$ in rms radians.
Similarly, the spectral density of the frequency fluctuations in Hz^2/Hz is given by

$$S_\nu(f) = \nu_0^2 \cdot S_y(f) = f^2 \cdot S_{\Delta\phi}(f) = K \cdot f^{x+2} \quad , \tag{42}$$

where ν_0 is the carrier frequency in hertz, and $S_y(f)$ is the spectral density of the fractional frequency fluctuations (see Section 6.1).

The integrated frequency jitter or residual FM is therefore

$$\Delta f^2 = \int_{f_1}^{f_2} S_\nu(f) \cdot df = \int_{f_1}^{f_2} K \cdot f^{x+2} \cdot df$$

$$\Delta f^2 = \frac{K}{x+3}(f_2^{x+3} - f_1^{x+3}) \text{ for } x \neq 3 \tag{43}$$

$$\Delta f^2 = K \cdot (\log f_2 - \log f_1) \text{ for } x = -3.$$

It is usually expressed as Δf in rms hertz.

The value of $S_\phi(f)$ in dB can be found from the more commonly used £(f) measure of SSB phase noise to carrier power ratio in dBc/Hz by adding 3 dB. The total integrated phase noise is obtained by summing the $\Delta\phi^2$ contributions from the straight-line approximations for each power law noise type. The ratio of total phase noise to signal power in the given integration bandwidth is equal to 10 log $\Delta\phi^2$.

References for Integrated Phase Jitter and Residual FM

1. W. J. Riley, "Integrate Phase Noise and Obtain Residual FM," *Microwaves*, August 1979, pp. 78-80.
2. U.L. Rohde, <u>Digital PLL Frequency Synthesizers</u>, pp. 411-418, Prentice-Hall, Englewood Cliffs, 1983.
3. D.A. Howe and T.N. Tasset, "Clock Jitter Estimation based on PM Noise Measurements," *Proc. 2003 Joint Mtg. IEEE Intl. Freq. Cont. Symp. and EFTF Conf.*, pp. 541-546, May 2003.

5.2.20.   Dynamic Stability

A dynamic stability analysis uses a sequence of sliding time windows to perform a dynamic Allan (DAVAR) or Hadamard (DHVAR) analysis, thereby showing changes (nonstationarity) in clock behavior versus time. It is able to detect variations in clock stability (noise bursts, changes in noise level or type, etc.) that would be difficult to see in an ordinary overall stability analysis. The results of a dynamic stability analysis are presented as a three-dimensional surface plot of log sigma versus log tau or averaging factor as a function of time or window number.

An example of a DAVAR plot is shown below. This example is similar to the one of Figure 2 in Reference [1], showing a source with white PM noise that changes by a factor of 2 at the middle of the record.
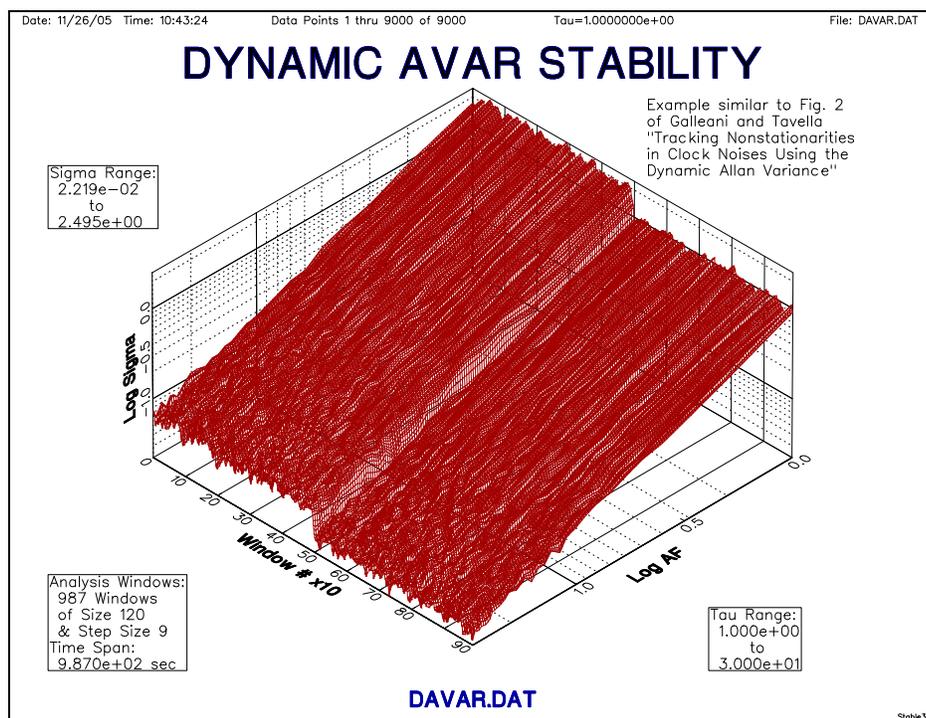
Figure 16. Example of a DAVAR plot.

References for Dynamic Stability

1. L. Galleani and P. Tavella, "Tracking Nonstationarities in Clock Noises Using the Dynamic Allan Variance," *Proc. 2005 Joint FCS/PTTI Meeting*.
2. L. Galleani and P. Tavella, "The Characterization of Clock Behavior with the Dynamic Allan Variance," *Proc. 2003 Joint FCS/EFTF Meeting*, pp. 239-244.

5.3. Confidence Intervals

It is wise to include error bars (confidence intervals) on a stability plot to indicate the degree of statistical confidence in the numerical results. The confidence limits of a variance estimate depend on the variance type, the number of data points and averaging factor, the statistical confidence factor desired, and the type of noise. This section describes the use of χ^2 statistics for setting the confidence intervals and error bars of a stability analysis.

It is generally insufficient to simply calculate a stability statistic such as the Allan deviation, thereby finding an estimate of its expected value. That determination should be accompanied by an indication of the confidence in its value as expressed by the upper and (possibly) lower limits of the statistic with a certain confidence factor. For example, if the estimated value of the Allan deviation is 1.0×10^{-11}, depending on the noise type and size of the data set, we could state with 95 % confidence that the actual value does not exceed (say) 1.2×10^{-11}. It is always a good idea to include such a confidence limit in reporting a statistical result, which can be shown as an upper numeric limit, upper and lower numeric bounds, or (equivalently) error bars on a plot. Even though those confidence limits or error bars are themselves inexact, they should be included to indicate the validity of the reported result.

If you are unfamiliar with the basics of confidence limits, it is recommended that an introductory statistics book be consulted for an introduction to this subject. For frequency stability analysis, the emphasis is on various variances, whose confidence limits (variances of variances) are treated with chi-squared (χ^2) statistics. Strictly speaking, χ^2 statistics apply to the classical standard variance, but they have been found applicable to all of the other variances (Allan, Hadamard, total, Thêo1, etc.) used for frequency stability analysis. A good introduction to confidence limits and error bars for the Allan variance may be found in Reference [1]. The basic idea is to (1) choose an single or double-sided confidence limits (upper or upper and lower bounds), (2) choose an appropriate confidence factor (e.g., 95 %), (3) determine the number of equivalent χ^2 degrees of freedom (edf), (4) use the inverse χ^2 distribution to find the normalized confidence limit(s), and (5) multiply those by the nominal deviation value to find the error bar(s).

5.3.1. Simple Confidence Intervals

The simplest confidence interval approximation, with no consideration of the noise type, sets the $\pm 1\sigma$ (68 %) error bars at $\pm \sigma_y(\tau)/\sqrt{N}$, where N is the number of frequency data points used to calculate the Allan deviation.

A more accurate determination of this confidence interval can be made by considering the noise type, which can be estimated by the B1 bias function (the ratio of the standard variance to the Allan variance). That noise type is then be used to determine a multiplicative factor, Kn, to apply to the confidence interval:

Noise Type	Kn
Random Walk FM	0.75
Flicker FM	0.77
White FM	0.87
Flicker PM	0.99
White PM	0.99

5.3.2   Chi-Squared Confidence Intervals

Chi-squared statistics can be applied to calculate single and double-sided confidence intervals at any desired confidence factor. These calculations are based on a determination of the number of degrees of freedom for the estimated noise type. Most stability plots show ±1σ error bars for its overlapping Allan deviation plot.

The error bars for the modified Allan and time variances are also determined by Chi-squared statistics, using the number of MVAR degrees of freedom for the particular noise type, averaging factor, and number of data points. During the Run function, noise type estimates are made at each averaging factor (except the last, where the noise type of the previous averaging factor is used).

Sample variances are distributed according to the expression

$$\chi^2 = \frac{edf \cdot s^2}{\sigma^2},$$
(44)

where χ^2 is the Chi-square, s^2 is the sample variance, σ^2 is the true variance, and edf is the equivalent number of degrees of freedom (not necessarily an integer). The edf is determined by the number of analysis points and the noise type. Procedures exist for establishing single- or double-sided confidence intervals with a selectable confidence factor, based on χ^2 statistics, for many of its variance functions. The general procedure is to choose a single- or double-limited confidence factor, p, calculate the corresponding χ^2 value, determine the edf from the variance type, noise type and number of analysis points, and thereby set the statistical limit(s) on the variance. For double-sided limits,

$$\sigma^2_{\min} = s^2 \cdot \frac{edf}{\chi^2(p, edf)} \text{ and } \sigma^2_{\max} = s^2 \cdot \frac{edf}{\chi^2(1-p, edf)}.$$
(45)

5.4.   Degrees of Freedom

The equivalent number of χ^2 degrees of freedom (edf) associated with a statistical variance (or deviation) estimate depends on the variance type, the number of data points, and the type of noise involved. In general, the progression from the original two-sample (Allan) variance to the overlapping, total, and Thêo1 variances has provided larger edfs and better confidence. The noise type matters because it determines the extent that the points are correlated. Highly correlated data have a smaller edf than does the same number of points of uncorrelated (white) noise. An edf determination therefore involves (1) choosing the appropriate algorithm for the particular variance type, (2) determining the dominant power law noise type of the data, and (3) using the number of data points to calculate the corresponding edf.

5.4.1.   AVAR, MVAR, TVAR, and HVAR EDF

The equivalent number of χ^2 degrees of freedom (edf) for the Allan variance (AVAR), the modified Allan variance (MVAR) and the related time variance (TVAR), and the Hadamard variance (HVAR) is found by a combined algorithm developed by C.A. Greenhall, based on its generalized autocovariance function [2].
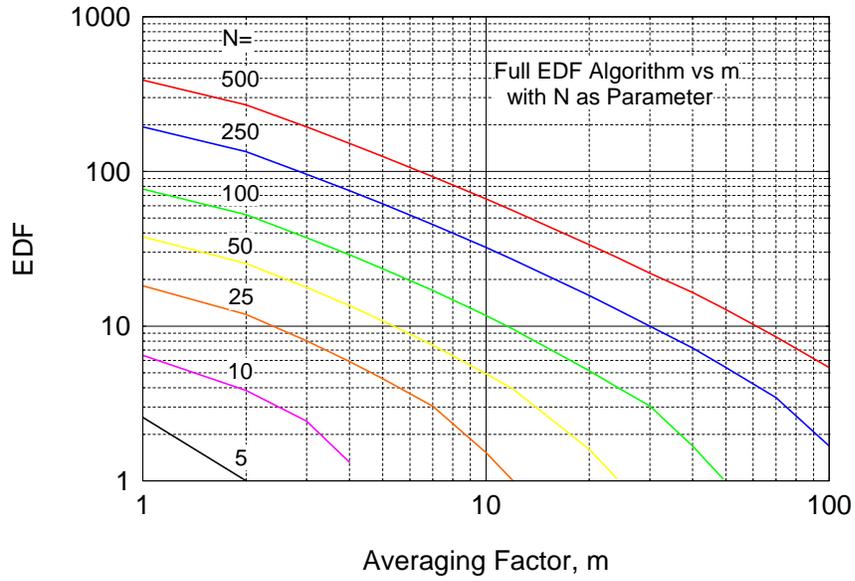
Figure 17. Overlapping ADEV EDF for W FM Noise.
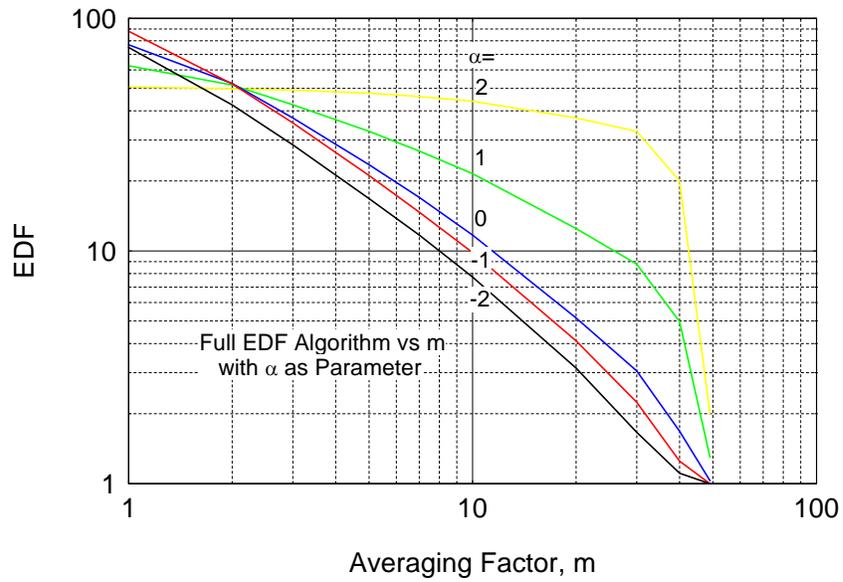
Figure 18. Overlapping ADEV EDF for N=100.
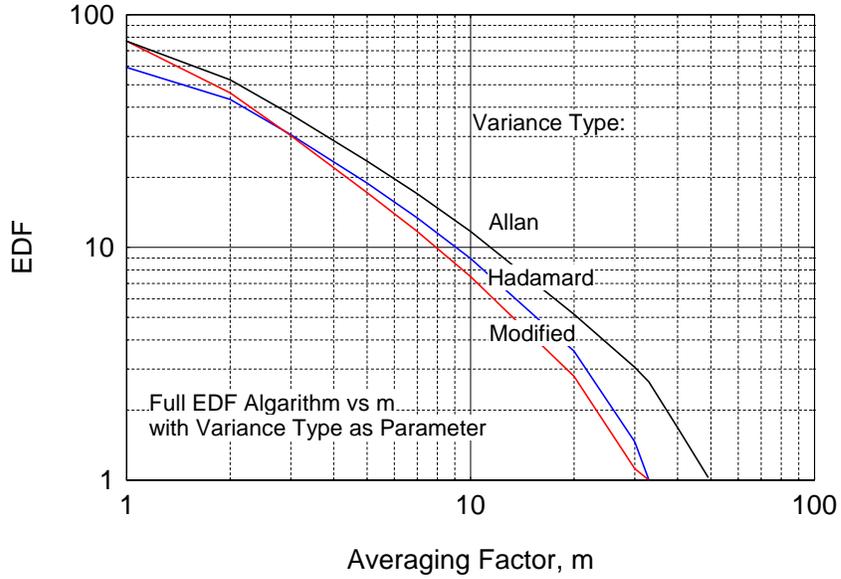
Figure 19. Overlapping EDF for N-100 and W FM Noise

This method for estimating the edf for the Allan, modified Allan, and Hadamard variances supersedes the following somewhat simpler empirical approximations (which may still be used).

The equivalent number of χ^2 degrees of freedom (edf) for the fully overlapping Allan variance (AVAR) can be estimated by the following approximation formulae for each power law noise type:

Table 5. AVAR approximation formulae for each power law noise type.

Power law noise type	AVAR edf, where N = # phase data points, m = averaging factor $= \tau/\tau_0$	
W PM	$$\frac{(N+1)(N-2m)}{2(N-m)}$$	
F PM	$$\exp\left[\ln\left(\frac{(N-1)}{2m}\right)\ln\left(\frac{(2m+1)(N-1)}{4}\right)\right]^{1/2}$$	
W FM	$$\left[\frac{3(N-1)}{2m} - \frac{2(N-2)}{N}\right]\frac{4m^2}{4m^2+5}$$	
F FM	$$\frac{2(N-2)^2}{2.3N-4.9}$$	For $m = 1$
	$$\frac{5N^2}{4m(N+3m)}$$	For $m > 1$
RW FM	$$\left[\frac{(N-2)}{m}\right]\left[\frac{(N-1)^2 - 3m(N-1) + 4m^2}{(N-3)^2}\right]$$	

The edf for the modified Allan variance (MVAR) can be estimated by the same expression as the overlapping Hadamard variance (see below) with the arguments changed as follows (valid for −2 ≤ α ≤ 2): MVAR and TVAR edf for N, m and α = MVAR edf for N + 1, m and α − 2.

The edf for the fully overlapping Hadamard variance (HVAR) can be found by an earlier algorithm also developed by C.A. Greenhall based on its generalized autocovariance function. The HVAR edf is found either as a summation (for small m cases with a small number of terms) or from a limiting form for large m, where 1/edf = (1/p)(a0−a1/p), with the coefficients as follows:

Table 6. HVAR edf coefficients.

Power law noise type	HVAR edf coefficients	
	a0	a1
W FM	7/9	1/2
F FM	1.00	0.62
RW FM	31/30	17/28
FW FM	1.06	0.53
RR FM	1.30	0.54

5.4.2.   TOTVAR EDF

The edf for the total variance (TOTVAR) is given by the formula $b(T/\tau) − c$, where T is the length of the data record, τ is the averaging time, and b and c are coefficients that depend on the noise type, as shown in Table 7:

Table 7. TOTVAR edf coefficients

Power law noise type	TOTVAR edf Coefficients	
	b	c
White FM	1.50	0
Flicker FM	1.17	0.22
Random walk FM	0.93	0.36

5.4.3.   MTOT EDF

The edf for the modified total variance (MTOT) is given by the same formula $b(T/\tau) − c$, where T is the length of the data record, τ is the averaging time, and b and c are coefficients that depend on the noise type as shown in Table 8:

Table 8. MTOT edf coefficients.

Power law noise type	MTOT edf Coefficients	
	b	c
White PM	1.90	2.10
Flicker PM	1.20	1.40
White FM	1.10	1.20
Flicker FM	0.85	0.50
Random walk FM	0.75	0.31

5.4.4. Thêo1 / ThêoH EDF

The equivalent number of χ^2 degrees of freedom (edf) for the Thêo1, hence, ThêoBR and ThêoH variances, is determined by the following approximation formulae for each power low noise type.

Table 9. ThêoBR and ThêoH approximation formulae for each power law noise type.

Power law noise type	Thêo1 edf, where N = # phase data points, $\tau = 0.75m$, m = averaging factor = τ/τ_0
White PM	$edf = \left(\dfrac{0.86(N_x + 1)(N_x - \frac{4}{3} \cdot \tau)}{N_x - \tau} \right) \left(\dfrac{\tau}{\tau + 1.14} \right)$
Flicker PM	$edf = \left(\dfrac{4.798 N_x^2 - 6.374 N_x \tau + 12.387 \tau}{(\tau + 36.6)^{1/2}(N_x - \tau)} \right) \left(\dfrac{\tau}{\tau + 0.3} \right)$
White FM	$edf = \left[\dfrac{4.1 N_x + 0.8}{\tau} - \dfrac{3.1 N_x + 6.5}{N_x} \right] \left(\dfrac{\tau^{3/2}}{\tau^{3/2} + 5.2} \right)$
Flicker FM	$edf = \left(\dfrac{2 N_x^2 - 1.3 N_x \tau - 3.5 \tau}{N_x \tau} \right) \left(\dfrac{\tau^3}{\tau^3 + 2.3} \right)$
Random walk FM	$edf = \left(\dfrac{4.4 N_x - 2}{2.9 \tau} \right) \left(\dfrac{(4.4 N_x - 1)^2 - 8.6\tau(4.4 N_x - 1) + 11.4\tau^2}{(4.4 N_x - 3)^2} \right)$

References for Confidence Intervals and Degrees of Freedom

1. D.A. Howe, D.W. Allan and J.A. Barnes, "Properties of Signal Sources and Measurement Methods," *Proc. 35th Annu. Symp. on Freq. Contrl.*, pp. 1-47, May 1981. Available on line at http://tf.nist.gov/general/publications.htm.
2. C. Greenhall and W. Riley, "Uncertainty of Stability Variances Based on Finite Differences," *Proc. 2003 PTTI Meeting*, December 2003.
3. D.A. Howe and T.K. Peppler, "Estimation of Very Long-Term Frequency Stability Using a Special-Purpose Statistic," *Proc. 2003 Joint Meeting of the European Freq. and Time Forum and the IEEE International Freq. Contrl. Symp.*, May 2003.
4. K. Yoshimura, "Degrees of Freedom of the Estimate of the Two-Sample Variance in the Continuous Sampling Method," *IEEE Transactions* IM-38, pp. 1044-1049, 1989.
5. "IEEE Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology - Random Instabilities," *IEEE Std 1139-1999*, July 1999.
6. C.A. Greenhall, "Estimating the Modified Allan Variance," *Proc. IEEE 1995 Freq. Contrl. Symp.*, pp. 346-353, May 1995.
7. D.A. Howe & C.A. Greenhall, "Total Variance: A Progress Report on a New Frequency Stability Characterization," *Proc. 1997 PTTI Meeting, pp. 39-48,* December 1997.
8. C.A. Greenhall, private communication, May 1999.
9. D.A. Howe, private communication, March 2000.
10. D.A. Howe, "Total Variance Explained," *Proc. 1999 Joint Meeting of the European Freq. and Time Forum and the IEEE Freq. Contrl. Symp.*, pp. 1093-1099, April 1999.
11. D.A. Howe, private communication, March 2000.
12. C.A. Greenhall, "Recipes for Degrees of Freedom of Frequency Stability Estimators," *IEEE Trans. Instrum. Meas.*, Vol. 40, No. 6, pp. 994-999, December 1991.
13. D.A. Howe, "Methods of Improving the Estimation of Long-Term Frequency Variance," *Proc. 11th European Freq. and Time Forum*, pp. 91-99, March 1997.
14. J.A. Barnes and D.W. Allan, "Variances Based on Data with Dead Time Between the Measurements," *NIST Technical Note 1318*, 1990.
15. D.A. Howe, R.L. Beard, C.A. Greenhall, F. Vernotte, W.J. Riley, and T.K. Peppler, "Enhancements to GPS Operations and Clock Evaluations Using a Total Hadamard Deviation," *IEEE T. Ultrason. Ferr.*, 52, pp. 1253-1261, Aug 2005..

5.5. Noise Identification

Identification of the dominant power law noise type is often necessary for setting confidence intervals and making bias corrections during a frequency stability analysis. The most effective means for power noise identification are based on the B_1 and $R(n)$ functions and the lag 1 autocorrelation.

5.5.1. Power Law Noise Identification

It is often necessary to identify the dominant power law noise process (WPM, FPM, WFM, FFM, RWFM, FWFM or RRFM) of the spectral density of the fractional frequency fluctuations, $S_y(f) = h_\alpha f^\alpha$ (α = 2 to –4), to perform a frequency stability analysis. For example, knowledge of the noise type is necessary to determine the equivalent number of chi-squared degrees of freedom (edf) for setting confidence intervals and error bars, and it is essential to know the dominant noise type to correct for bias in the newer Total and Thêo1 variances. While the noise type may be known a priori or estimated manually, it is desirable to have an analytic method for power law noise identification that can be used automatically as part of a stability analysis algorithm.

There is little literature on the subject of power-law noise identification. The most common method for power law noise identification is simply to observe the slope of a log-log plot of the Allan or modified Allan deviation versus averaging time, either manually or by fitting a line to it. This obviously requires at least two stability points. During a stability calculation, it is desirable (or necessary) to automatically identify the power law noise type at each point, particularly if bias corrections and/or error bars must be applied.

5.5.2. Noise Identification Using B₁ and R(n)

A noise identification algorithm that has been found effective in actual practice, and that works for a single τ point over the full range of $-4 \leq \alpha \leq 2$ is based on the Barnes B_1 function, which is the ratio of the N-sample (standard) variance to the two-sample (Allan) variance, and the R(n) function [1], which is the ratio of the modified Allan to the normal Allan variances. The B_1 function has as arguments the number of frequency data points, N, the dead time ratio, r (which is set to 1), and the power law t-domain exponent, μ. The B_1 dependence on μ is used to determine the power law noise type for $-2 \leq \mu \leq 2$ (W and F PM to FW FM). For a B_1 corresponding to $\mu = -2$, the $\alpha = 1$ or 2 (F PM or W PM noise) ambiguity can be resolved with the R(n) ratio using the modified Allan variance. For the Hadamard variance, for which RR FM noise can apply, (m =3, a = –4), the B_1 ratio can be applied to frequency (rather than phase) data, and adding 2 to the resulting μ.

The overall noise B_1/R(n) noise identification process is therefore:

1. Calculate the standard and Allan variances for the applicable τ averaging factor.
2. Calculate $B_1(N, r=1, \mu) = \dfrac{N(1 - N^{\mu})}{2(N - 1)(1 - 2^{\mu})}$.
3. Determine the expected B_1 ratios for a = –3 through 1 or 2.
4. Set boundaries between them and find the best power law noise match.
5. Resolve an $\alpha = 1$ or 2 ambiguity with the modified Allan variance and R(n).
6. Resolve an $\alpha = -3$ or –4 ambiguity by applying B_1 to frequency data.

The boundaries between the noise types are generally set as the geometric means of their expected values. This method cannot distinguish between W and F PM at unity averaging factor.

5.5.3. The Autocorrelation Function
The autocorrelation function (ACF) is a fundamental way to describe a time series by multiplying it by a delayed version of itself, thereby showing the degree by which its value at one time is similar to its value at a certain later time. More specifically, the autocorrelation at lag k is defined as

$$\rho_k = \frac{E[(z_t - \mu)(z_{t+k} - \mu)]}{\sigma_z^2} , \tag{46}$$

where z_t is the time series, μ is its mean value, σ_z^2 is its variance, and E denotes the expected value. The autocorrelation is usually estimated by the expression

$$r_k = \frac{\dfrac{1}{N}\sum_{t=1}^{N-k}(z_t - \bar{z})(z_{t+k} - \bar{z})}{\dfrac{1}{N}\sum_{t=1}^{N}(z_t - \bar{z})^2} , \tag{47}$$

where $\bar{z}$ is the mean value of the time series and N is the number of data points [2].

5.5.4. The Lag 1 Autocorrelation

The lag 1 autocorrelation is simply the value of r_1 as given by the expression above. For frequency data, the lag 1 autocorrelation is able to easily identify white and flicker PM noise, and white (uncorrelated) FM noise, for which the expected values are –1/2, –1/3 and zero, respectively. The more divergent noises have positive r_1 values that depend on the number of samples, and tend to be larger (approaching 1). For those more divergent noises, the data are differenced until they become stationary, and the same criteria as for WPM, FPM and WFM are then used, corrected for the differencing. The results can be rounded to determine the dominant noise type or used directly to estimate the noise mixture.

5.5.5. Noise Identification Using r_1

An effective method for identifying power law noises using the lag 1 autocorrelation [3] is based on the properties of discrete-time fractionally integrated noises having spectral densities of the form $(2 \sin \pi f)^{-2\delta}$. For $\delta < \frac{1}{2}$, the process is stationary and has a lag 1 autocorrelation equal to $\rho_1 = \delta / (1-\delta)$ [4], and the noise type can therefore be estimated from $\delta = r_1 / (1+r_1)$. For frequency data, white PM noise has $\rho_1 = -1/2$, flicker PM noise has $\rho_1 = -1/3$, and white FM noise has $\rho_1 = 0$. For the more divergent noises, first differences of the data are taken until a stationary process is obtained as determined by the criterion $\delta < 0.25$. The noise identification method therefore uses $p = -\text{round}(2\delta) -2d$, where round (2δ) is 2δ rounded to the nearest integer and d is the number of times that the data is differenced to bring δ down to < 0.25. If z is a τ-average of frequency data y(t), then $\alpha = p$; if z is a τ-sample of phase data x(t), then $\alpha = p + 2$, where α is the usual power law exponent f^α, thereby determining the noise type at that averaging time. The properties of this power law noise identification method are summarized in Table 10. It has excellent discrimination for all common power law noises for both phase and frequency data, including difficult cases with mixed noises.

Noise Type	α	Phase Data* x(t)	d=0 ACF of Phase Data
W PM	2		
F PM	1		
W FM	0		
F FM	–1		
RW FM	–2		
* The differencing operation changes the appearance of the phase data to that shown 2 rows higher.			

Figure 20. Lag 1 Autocorrelation for Various Power Law Noises

Table 10. Lag 1 Autocorrelation for Various Power Law Noises and Differences

Noise Type α	Lag 1 Autocorrelation, r₁[†]					
	d=0		d=1		d=2	
	x(t)	y(t)	x(t)	y(t)	x(t)	y(t)
2	0	−1/2	−1/2	−2/3	−2/3	−3/4
1	≈0.7	−1/3	−1/3	−3/5	−3/5	−5/7
0	≈1	0	0	−1/2	−1/2	−2/3
-1	≈1	≈0.7	≈0.7	−1/3	−1/3	−3/5
-2	≈1	≈1	≈1	0	0	−1/2
[†] Shaded values are those used for noise ID for the particular noise and data type.						

5.5.6. Noise ID Algorithm

The basic lag 1 autocorrelation power law noise identification algorithm is quite simple. The inputs are a vector $z_1,\ldots,$ z_N of phase or frequency data, the minimum order of differencing dmin (default = 0), and the maximum order of differencing dmax. The output is p, an estimate of the α of the dominant power law noise type, and (optionally) the value of d.

Done = False, $d = 0$
While Not Done

$$\bar{z} = \frac{1}{N}\sum_{i=1}^{N} z_i$$

$$r_1 = \frac{\sum_{i=1}^{N-1}(z_i - \bar{z})(z_{i+1} - \bar{z})}{\sum_{i=1}^{N}(z_i - \bar{z})^2}$$

$$\delta = \frac{r_1}{1 + r_1}$$

If d >= dmin And ($\delta < 0.25$ Or d >= dmax)

$\quad\quad p = -2(\delta + d)$

$\quad\quad$ Done = True

Else

$\quad\quad\quad z_1 = z_2 - z_1, \ldots, z_{N-1} = z_N - z_{N-1}$

$\quad\quad\quad N = N - 1$

$\quad\quad\quad d = d + 1$

End If
End While
Note: May round p to nearest integer

Figure 21. The basic lag 1 autocorrelation power law noise identification algorithm.

The input data should be for the particular averaging time, τ, of interest, and it may therefore be necessary to decimate the phase data or average the frequency by the appropriate averaging factor before applying the noise identification algorithm. The dmax parameter should be set to two or three for an Allan or Hadamard (two or three-sample) variance analysis, respectively. The alpha result is equal to p+2 or p for phase or frequency data, respectively, and may be rounded to an integer (although the fractional part is useful for estimated mixed noises). The algorithm is fast, requiring only the calculation of one autocorrelation value and first differences for several times. It is independent of any particular variance. The lag 1 autocorrelation method yields good results, consistently identifying pure power noise for $\alpha = 2$ to –4 for sample sizes of about 30 or more, and generally identifying the dominant type of mixed noises when it is at least 10 % larger than the others. For a mixture of adjacent noises, the fractional result provides an indication of their ratio. It can handle all averaging factors.

Before analysis, the data should be preprocessed to remove outliers, discontinuities, and deterministic components. Acceptable results can be obtained from the lag 1 autocorrelation noise identification method for $N \geq 32$, where N is the number of data points. The algorithm tends to produce jumps in the estimated alpha for mixed noises when the differencing factor, d, changes (although the alpha value when rounded to an integer is still consistent). This can be avoided by using the same d for the entire range of averaging times, at the expense of higher variability when a lower d would have been sufficient. The lag 1 autocorrelation method for power law noise identification is a fast and effective way to support the setting of confidence intervals and to apply bias corrections during a frequency stability analysis, as shown Figure 22:
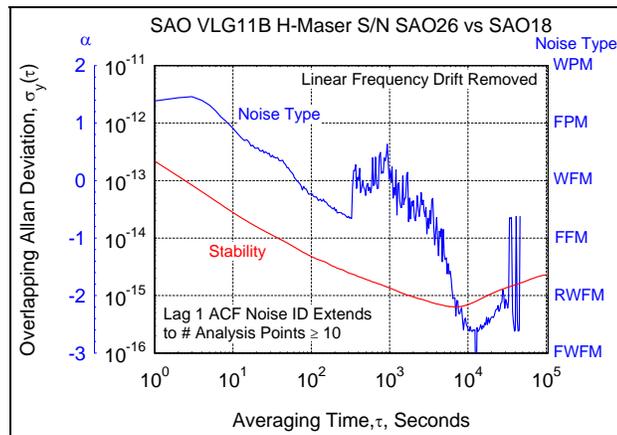
Figure 22. Frequency Stability and Noise Analysis of Two Hydrogen Masers

References for Noise Identification

1. D.A. Howe, R.L. Beard, C.A. Greenhall, F. Vernotte, W.J. Riley and T.K. Peppler, "Enhancements to GPS Operations and Clock Evaluations Using a "Total" Hadamard Deviation," *IEEE Trans. Ultrasonics, Ferrelectrics and Freq. Contrl.*, Vol. 52, No. 8, pp. 1253-1261, Aug. 2005.
2. J.A. Barnes, "Tables of Bias Functions, B1 and B2, for Variances Based on Finite Samples of Processes with Power Law Spectral Densities," *NBS Technical Note 375*, Jan. 1969.
3. D.B Sullivan, D.W. Allan, D.A. Howe and F.L. Walls (editors), *Characterization of Clocks and Oscillators*, NIST Technical Note 1337, Sec. A-6, 1990.
4. G. Box and G. Jenkins, Time Series Analysis, Forecasting and Control, Chapter 2, Holden-Day, Oakland, California, 1976, ISBN 0-8162-1104-3.
5. W.J. Riley and C.A. Greenhall, "Power Law Noise Identification Using the Lag 1 Autocorrelation," *Proc. 18th European Frequency and Time Forum*, April 2004.
6. P. Brockwell and R. Davis, Time Series: Theory and Methods, 2nd Edition, Eq. (13.2.9), Springer-Verlag, New York, 1991.
7. N.J. Kasdin and T. Walter, "Discrete Simulation of Power Law Noise," *Proc. 1992 IEEE Frequency Control Symposium*, pp. 274-283, May 1992.
8. J. McGee and D.A. Howe, "ThêoH and Allan Deviation as Power-Law Noise Estimators," *IEEE Trans. Ultrasonics, Ferroelectrics and Freq. Contrl.*, Feb. 2007

5.6. Bias Functions

Several bias functions are defined and used in the analysis of frequency stability, as defined below. In particular, B_1, the ratio of the standard variance to the Allan variance, and R(n), the ratio of the modified Allan variance to the normal Allan variance, are used for the identification of power law noise types (see section 5.2.2), and the B_2 and B_3 bias functions are used to correct for dead time in a frequency stability measurement.

5.7. B_1 Bias Function

The B_1 bias function is the ratio of the N-sample (standard) variance to the two-sample (Allan) variance with dead time ratio r = T/τ, where T = time between measurements, τ = averaging time, and μ = exponent of τ in Allan variance for a certain power law noise process:

$$B_1(N,r,\mu) = \frac{\sigma^2(N,T,\tau)}{\sigma^2(2,T,\tau)}. \tag{48}$$

The B_1 bias function is useful for performing power law noise identification by comparing its actual value to those expected for the various noise types (see section 5.2.2).

5.8. B_2 Bias Function

The B_2 bias function is the ratio of the two-sample (Allan) variance with dead time ratio r = T/τ to the two-sample (Allan) variance without dead time (r = 1):

$$B_2(r,\mu) = \frac{\sigma^2(2,T,\tau)}{\sigma^2(2,\tau,\tau)}. \tag{49}$$

5.9. B₃ Bias Function

The B_3 bias function is the ratio of the N-sample (standard) variance with dead time ratio r = T/τ at multiples M = τ/τ₀ of the basic averaging time τ_0 to the N-sample variance with the same dead time ratio at averaging time τ:

$$B_3(N,M,r,\mu) = \frac{\sigma^2(N,M,T,\tau)}{\sigma^2(N,T,\tau)} \qquad . \tag{50}$$

The product of the B_2 and B_3 bias functions is used for dead time correction, as discussed in section 5.7.

5.10. R(n) Bias Function

The R(n) function is the ratio of the modified Allan variance to the normal Allan variance for n = number of phase data points. Note: R(n) is also a function of α, the exponent of the power law noise type:

$$R(n) = \frac{\text{Mod } \sigma_y^2(\tau)}{\sigma_y^2(\tau)} . \tag{51}$$

The R(n) bias function is useful for performing power law noise identification by comparing its actual value to those expected for the various noise types (see Section 5.2.2).

5.11. TOTVAR Bias Function

The TOTVAR statistic is an unbiased estimator of the Allan variance for white and flicker PM noise, and for white FM noise. For flicker and random walk FM noise, TOTVAR is biased low as τ becomes significant compared with the record length. The ratio of the expected value of TOTVAR to AVAR is given by the expression

$$B(\text{TOTAL}) = 1 - a \cdot \left(\frac{\tau}{T}\right), \quad 0 < \tau < \frac{T}{2}, \tag{52}$$

where a = 1/3·ln2 = 0.481 for flicker FM noise, a = 3/4 = 0.750 for random walk FM noise, and T is the record length. At the maximum allowable value of τ = T/2, TOTVAR is biased low by about 24 % for RW FM noise. This bias function should be used to correct all reported TOTVAR results.

5.12. MTOT Bias Function

The MTOT statistic is a biased estimator of the modified Allan variance. The MTOT bias factor (the ratio of the expected value of Mod Totvar to MVAR), as shown in Table 11, depends on the noise type but is essentially independent of the averaging factor and number of data points.

Table 11. MTOT bias factors for each noise type.

Noise	Bias Factor
W PM	1.06
F PM	1.17
W FM	1.27
F FM	1.30
RW FM	1.31

This bias factor should be used to correct all reported MTOT results.

5.13. Thêo1 Bias

The Thêo1 statistic is a biased estimator of the Allan variance. The Thêo1 bias factor (the ratio of the expected value of Thêo1 to AVAR) depends on both noise type and averaging factor:

$$\text{Thêo1 Bias} = \frac{\text{AVAR}}{\text{Thêo1}} = a + \frac{b}{m^c} \, , \tag{53}$$

where m is the averaging factor and the constants a, b and c are given in Table 12. Note that the effective tau for a Thêo1 estimation is $t = 0.75 \cdot m \cdot t_0$, where t_0 is the measurement interval.

Table 12. Constants a, b, and c for Thêo1 bias.

Noise	Alpha	a	b	c
RW FM	–2	2.70	–1.53	0.85
F FM	–1	1.87	–1.05	0.79
W FM	0	1.00	0.00	0.00
F PM	1	0.14	0.82	0.30
W PM	2	0.09	0.74	0.40

5.14. ThêoH Bias

ThêoH statistic is a bias-removed estimator of the Allan variance. ThêoH is the best statistic for estimating the frequency stability and type of noise over the widest possible range of averaging times without explicit knowledge of the noise type.

References for Bias Functions

1. "IEEE Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology - Random Instabilities," *IEEE Std 1139-1999*, July 1999.
2. C.A. Greenhall, "Estimating the Modified Allan Variance," *Proc. IEEE 1995 Freq. Contrl. Symp.*, pp. 346-353, May 1995.
3. D.A. Howe & C.A. Greenhall, "Total Variance: A Progress Report on a New Frequency Stability Characterization," *Proc. 1997 PTTI Meeting, pp. 39-48,* December 1997.
4. C.A. Greenhall, private communication, May 1999.
5. D.A. Howe, private communication, March 2000.
6. D.A. Howe, "Total Variance Explained," *Proc. 1999 Joint Meeting of the European Freq. and Time Forum and the IEEE Freq. Contrl. Symp.*, pp. 1093-1099, April 1999.
7. C.A. Greenhall, "Recipes for Degrees of Freedom of Frequency Stability Estimators," *IEEE Trans. Instrum. Meas.*, Vol. 40, No. 6, pp. 994-999, December 1991.
8. D.A. Howe, "Methods of Improving the Estimation of Long-Term Frequency Variance," *Proc. 11th European Freq. and Time Forum*, pp. 91-99, March 1997.
9. J.A. Barnes and D.W. Allan, "Variances Based on Data with Dead Time Between the Measurements," *NIST Technical Note 1318*, 1990.
10. C.A Greenhall, private communication, May 2000.
11. D.A. Howe, R.L. Beard, C.A. Greenhall, F. Vernotte, W.J. Riley and T.K. Peppler, "Enhancements to GPS Operations and Clock Evaluations Using a "Total" Hadamard Deviation," *IEEE Trans. Ultrasonics, Ferrelectrics and Freq. Contrl.*, Vol. 52, No. 8, pp. 1253-1261, Aug. 2005.

5.15. Dead Time

Dead time can occur in frequency measurements because of instrumentation delay between successive measurements, or because of a deliberate wait between measurements. It can have a significant effect on the results of a stability analysis, especially for the case of large dead time (e.g., frequency data taken for 100 seconds, once per hour).

> Dead time can occur in frequency measurements and can significantly affect a subsequent stability analysis. Methods are available to correct for dead time and thus obtain unbiased results.
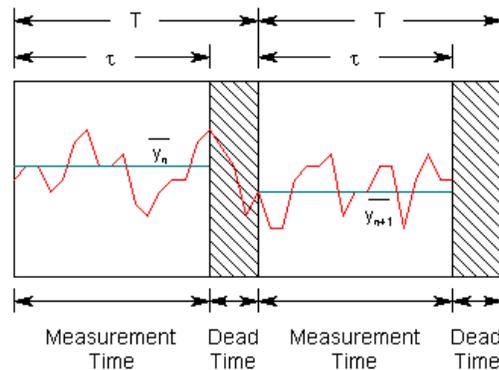
Figure 23. Illustration of dead time between successive measurements.

Dead time corrections can be applied by dividing the calculated Allan deviation by the square root of the product of the Barnes B_2 and B_3 bias ratios. These corrections are particularly important for non-white FM noise with a large dead time ratio. Restricting the dead time corrections to Allan deviations is a conservative approach based on the B_2 and B_3 definitions. Those bias functions depend critically on the power law noise type. Requiring manual noise selection avoids the problem of noise identification for biased data having the wrong sigma-tau slope. Dead time

correction is problematic for data having multiple noise types. In addition to introducing bias, measurement dead time reduces the confidence in the results, lowers the maximum allowable averaging factor, and prevents proper conversion of frequency to phase. Moreover, no information is available about the behavior of the device under test during the dead time. It is recommended that these issues be avoided by making measurements with zero dead time.

Dead time that occurs at the end of a measurement can be adjusted for in an Allan deviation determination by using the Barnes B_2 bias function [1], the ratio of the two-sample variance with dead time ratio $r = T/\tau$ to the two-sample variance without dead time. Otherwise, without this correction, one can determine only the average frequency and its drift. When such data are used to form frequency averages at longer tau, it is necessary to also use the B_3 bias function [2], the ratio of the variance with distributed dead time to the variance with all the dead time at the end. Those bias corrections are made by use of the product of B_2 and B_3. The power law noise type must be known in order to calculate these bias functions. Simulated periodically sampled frequency data with distributed dead time for various power law noise processes shows good agreement with the B_2 and B_3 bias function corrections, as shown in Figure 24.
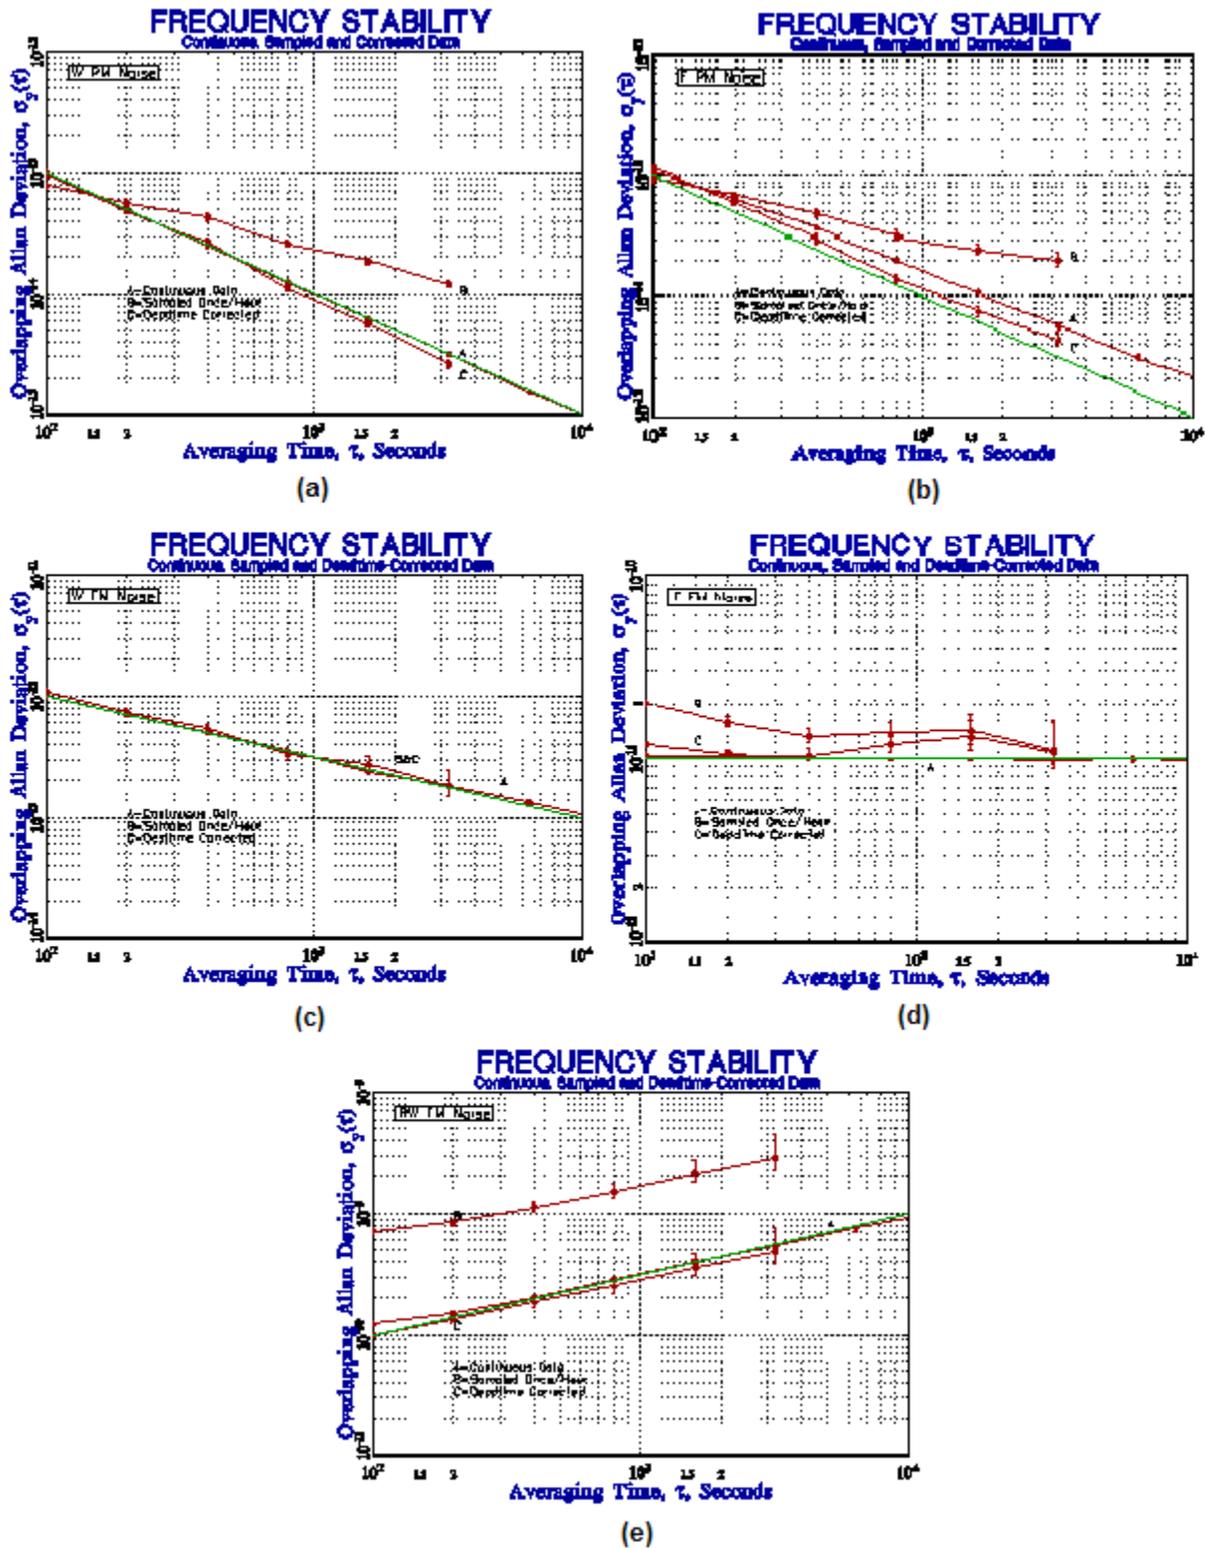
Figure 24. Frequency stability plots for common power law noises with large measurement dead time (r = T/τ = 36). Simulated data sampled for τ = 100 seconds once per hour for 10 days. Nominal 1×10^{-11} stability at τ = 100 seconds shown by lines. Plots show stability of simulated data sets for continuous, sampled and dead time-corrected data. (a) White PM (μ = -2) $\sqrt{B_2}$ = 0.82 at AF = 1, (b) Flicker PM (μ = -2), $\sqrt{B_2}$ = 0.82 at AF = 1, (c) White FM (μ = -1), $\sqrt{B_2}$ = 1.00 at AF = 1, (d) Flicker FM (μ = 0) $\sqrt{B_2}$ = 1.92 at AF = 1, (e) RW FM (μ = 1) $\sqrt{B_2}$ = 7.31 at AF = 1.

These simulations show that the B_2 and B_3 bias corrections are able to support reasonably accurate results for sampled frequency stability data having a large dead time, when the power law noise type is known. The slope of the raw sampled stability plot does not readily identify the noise type, however, and mixed noise types would make correction difficult. The relatively small number of data points reduces the confidence of the results, and limits the allowable averaging factor. Moreover, the infrequent measurements provide no information about the behavior of the clock during the dead time, and prevent a proper conversion of frequency to phase. Sparsely sampled data are therefore not recommended for the purpose of stability analysis.

References for Dead Time

1. J.A. Barnes, "Tables of Bias Functions, B1 and B2, for Variances Based on Finite Samples of Processes with Power Law Spectral Densities," *NBS Technical Note 375*, January 1969.
2. J.A. Barnes and D.W. Allan, "Variances Based on Data with Dead Time Between the Measurements," *NIST Technical Note 1318*, 1990.

5.16. Unevenly Spaced Data

Unevenly spaced phase data can be handled if they have associated timetags by using the individual timetag spacing when converting them to frequency data. Then, if the tau differences are reasonably small, the data may be analyzed by use of the average timetag spacing as the analysis tau, in effect placing the frequency data on an average uniform grid. While completely random data spacing is not amenable to this process, tau variations of ±10 % will yield reasonable results as long as the exact interval is used for phase to frequency conversion.

An example of unevenly spaced data is two-way satellite time and frequency transfer (TWSTFT) measurements made on Monday, Wednesday, and Friday of each week, where the data spacing is either one or two days.
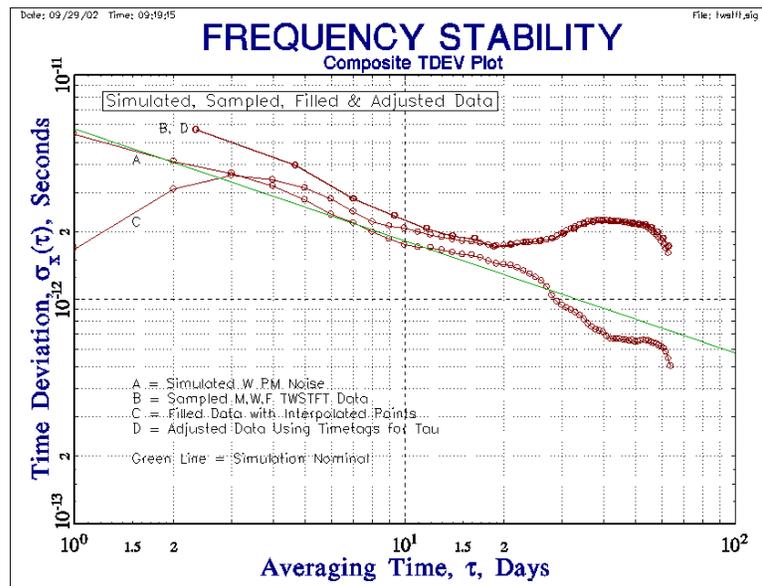
Figure 25. TDEV results for simulated TWSTFT data.

The TWSTFT data are simulated as 256 points of white PM noise with an Allan deviation (ADEV) level of $\sigma_y(\tau) = 1 \times 10^{-11}$ at 1-day. A composite plot of the TWSTFT TDEV results is shown above. The corresponding TDEV is 5.77×10^{-12} sec at $\tau = 1$ day (TDEV = MDEV divided by $\sqrt{3}$), as shown in curve A. Note that these time-stability plots

include points at all possible tau values. The green line shows that the –0.5 slope of the TDEV plot for W PM noise. The TWSTFT data are sampled once on Monday, Wednesday, and Friday of each week. These sampled data therefore have an average tau of 7/3 = 2.33 days, and their TDEV is shown in curve B. If the missing points are replaced by linearly interpolated phase values, the TDEV becomes highly distorted, as shown in curve C. If the sampled phase data are converted to frequency data using their timetag differences to determine the individual measurement intervals, the average tau, τ_{avg}, is close to 2.33 days (depending on the final fractional week that is included), and the resulting TDEV is shown in curve D. It is essentially identical to that for the sampled phase data shown in curve B. It is interesting that, although the converted frequency results are different depending on whether the average or individual (more correct) taus are used, the (integrated) TDEV results are not (at least for a white PM noise process).

None of the results is in good agreement with the nominal simulation. The result with the linearly interpolated phase points is particularly bad for $\tau<\tau_{avg}$, and is similar to that of Tavella and Leonardi, as shown in Figure 1 of Reference [1]. As they point out in that paper, because the true sampling interval is τ_{avg}, it is not possible to estimate the noise at shorter times, especially for an uncorrelated white noise process. They further suggest that the higher level of the estimated noise is related to the ratio of the true and interpolated sampling times (≈ 2.33) and the $\sqrt{\tau}$ dependence of TDEV. By applying a correction factor of $\sqrt{2.33} \approx 1.5$, the longer-tau TDEV estimates are lowered to the correct level. These factors are smaller for other non-white PM and FM noise processes. The adjusted method of using frequency data converted from phase data by using individual tau values adjusted for the timetag spacing is recommended because it does not use interpolation, does not present results at unrealistically low tau, and uses the best frequency estimates.

Another situation is data that are taken in bursts. In that case, the best approach is probably to analyze the segments separately, perhaps averaging those results to obtain better statistical confidence. One could obtain reasonable results for the shorter averaging times, but cannot apply standard techniques to analyze the complete data set.

References for Unevenly Spaced Data

1. P. Tavella and M. Leonardi, "Noise Characterization of Irregularly Spaced Data," *Proceedings of the 12th European Frequency and Time Forum*, pp. 209-214, March 1998.
2. C. Hackman and T.E. Parker, "Noise Analysis of Unevenly Spaced Time Series Data," *Metrologia*, Vol. 33, pp. 457-466, 1996.

5.17. Histograms

A histogram shows the amplitude distribution of the phase or frequency fluctuations, and can provide insight regarding them. We can expect a normal (Gaussian) distribution for a reasonably sized data set, and a different (e.g., bimodal) distribution can be a sign of a problem.

For a normal distribution, the standard deviation is approximately equal to the half-width at half-height (HWHA = 1.177s) .

Standardized Normal Distribution (μ=0, σ=1)

$$f(x)=(1/\sqrt{2\pi})e^{-(x^2/2)}$$

±1σ

Half-Amplitude

-1.177σ   +1.177σ

f(x)

x (In units of σ)
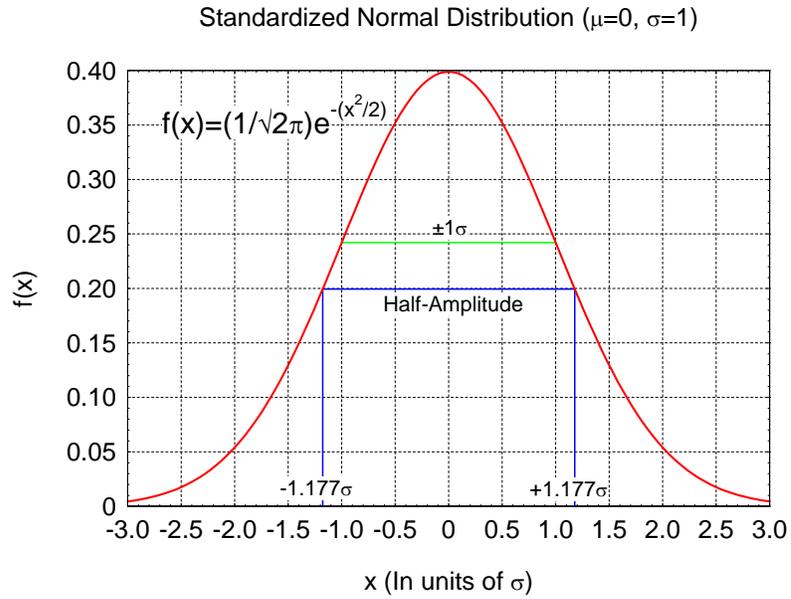
Figure 26. Half-width at half-height for the standard deviation.

HISTOGRAM
Simulated White FM Noise σ=1e-12

Data Units: pp10^{12}

\# Points Per Bin

Data Value

Normal Curve Fit: Mean=1.0161202e−02, Std Dev=1.0007094e+00, Peak=972
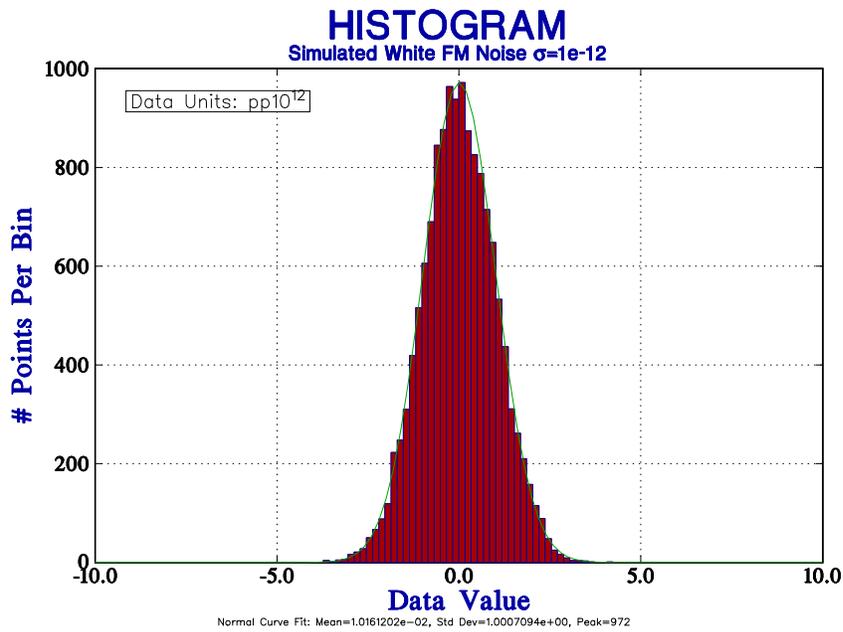
Figure 27. An example of a histogram for a set of white FM noise.

5.18. Frequency Offset

It is often necessary to estimate the frequency offset from either phase or frequency data.

Frequency offset is usually calculated from phase data by either of three methods:

1.   A least squares linear fit to the phase data (optimum for white PM noise):

$x(t) = a + bt$, where slope $= y(t) = b$.

2.   The average of the first differences of the phase data (optimum for white FM noise):

$y(t) =$ slope $= [x(t+\tau) – x(t)]/\tau$.

3.   The difference between the first and last points of the phase data:

$y(t) =$ slope $= [\ x(end) – x(start)\ ] / (M–1)$, where $M = $ # phase data points.

This method is used mainly to match the two endpoints.

5.19. Frequency Drift

Most frequency sources have frequency drift, and it is often necessary (and usually advisable) to characterize and remove this systematic drift before analyzing the stochastic noise of the source. The term drift refers to the systematic change in frequency due to all effects, while aging includes only those effects internal to the device. Frequency drift is generally analyzed by fitting the trend of the frequency record to an appropriate mathematical model (e.g., linear, log, etc.), often by the method of least squares. The model may be based on some physical basis or simply a convenient equation, using either phase or frequency data, and its suitability may be judged by the degree to which it produces white (i.e., uncorrelated) residuals.

Frequency drift is the systematic change in frequency due to all effects, while frequency aging is the change in frequency due to effects within the device. Thus, for a quartz crystal oscillator, aging refers to a change in the resonant frequency of its quartz crystal resonator, while drift would also include the effects of its external environment. Therefore, drift is the quantity that is usually measured, but it is generally done under constant environmental conditions to the greatest extent possible so as to approximate the aging of the device.

5.20. Drift Analysis Methods

Several drift methods are useful for phase or frequency data as described below. The best method depends on the quality of the fit, which can be judged by the randomness of the residuals.

Table 13. Drift analysis methods for phase or frequency data.

Data	Method	Noise Model
Phase	Quadratic Fit	W PM
Phase	Avg of 2nd Diffs	RW FM
Phase	3-Point Fit	W & RW FM
Phase	Linear Fit	Frequency Offset
Phase	Avg of 1st Diffs	Frequency Offset
Phase	Endpoints	Frequency Offset
Freq	Linear Fit	W FM
Freq	Bisection Fit	W & RW FM
Freq	Log Fit	Stabilization
Freq	Diffusion Fit	Diffusion

5.21. Phase Drift Analysis

Three methods are commonly used to analyze frequency drift in phase data:

1. A least squares quadratic fit to the phase data:

$x(t) = a + bt + ct^2$, where $y(t) = x'(t) = b + 2ct$, slope $= y'(t) = 2c$.

This continuous model can be expressed as $x_n = a + \tau_0 bn + \tau_0^2 cn^2$ for $n = 1, 2, 3..., N$, and τ_0 is the sampling interval for discrete data, where the a, b and c coefficients have units of sec, sec/sec and sec/sec^2, respectively, and the frequency drift slope and intercept are $2c$ and b, respectively. The fit coefficients can be estimated by the following expressions [1]:

$$\hat{a} = \left(A\sum_{n=1}^{N} x_n + \tau_0 B\sum_{n=1}^{N} nx_n + \tau_0^2 C\sum_{n=1}^{N} n^2 x_n \right) / G$$

$$\hat{b} = \left(B\sum_{n=1}^{N} x_n + \tau_0 D\sum_{n=1}^{N} nx_n + \tau_0^2 E\sum_{n=1}^{N} n^2 x_n \right) / G\tau_0 ,$$ (54)

$$\hat{c} = \left(C\sum_{n=1}^{N} x_n + \tau_0 E\sum_{n=1}^{N} nx_n + \tau_0^2 F\sum_{n=1}^{N} n^2 x_n \right) / G\tau_0^2$$

where the A-F terms are as follows:

$A = 3\big[3N(N+1)+2\big]$
$B = -18(2N+1)$
$C = 30$
$D = 12(2N+1)(8N+11)/\big[(N+1)(N+2)\big]$
$E = -180/(N+2)$
$F = 180/\big[(N+1)(N+2)\big]$
$G = N(N-1)(N-2)$

A quadratic fit to the phase data is the optimum model for white PM noise.

2. The average of the second differences of the phase data:

$y(t) = [x(t+\tau)-x(t)]/\tau$, slope $= [y(t+\tau)-y(t)]/\tau = [x(t+2\tau)-2x(t+\tau)+x(t)]/\tau^2$.

This method is optimum for random walk FM noise.

3. A three-point fit at the start, middle, and end of the phase data:

slope $= 4[x(end)-2x(mid)+x(start)]/(M\tau)^2$, where M = the number of data points.

It is the equivalent of the bisection method for frequency data.

5.22. Frequency Drift Analysis

Four methods are commonly used to analyze frequency drift in frequency data:

1. A least squares linear regression to the frequency data:

$y(t) = a + bt$, where a = intercept, b = slope $= y'(t)$.

Linear frequency drift can be estimated by a linear least squares fit to the frequency data, $y(t) = a + bt$. That continuous model can be expressed as $y_n = a + \tau_0 bn$ for $n = 1, 2, 3..., M$ where M is the number of frequency data points, τ_0 is the sampling interval for the discrete data. The frequency drift intercept and slope are a and b, and have units of sec and sec/sec, respectively. The fit coefficients can be estimated by equations 55 and 56:

$$\hat{b} = \frac{M\sum_{n=1}^{M} ny_n - \sum_{n=1}^{M} n \sum_{n=1}^{M} y_n}{M\sum_{n=1}^{M} n^2 - \left(\sum_{n=1}^{M} n\right)^n} \tag{55}$$

and

$$\hat{a} = \frac{\sum_{n=1}^{M} y_n}{M} - b \cdot \frac{\sum_{n=1}^{M} n}{M}. \tag{56}$$

A linear fit to the frequency data is the optimum model for white FM noise.

2. The frequency averages over the first and last halves of the data:

slope = 2 [y(2nd half) – y(1st half)] / (Nτ), where N = number of points.

This bisection method is optimum for white and random walk FM noise.

3. A log model of the form (see MIL-O-55310B) that applies to frequency stabilization:

$y(t) = a \cdot \ln(bt+1)$, where slope = $y'(t) = ab/(bt+1)$.

4. A diffusion ($\sqrt{t}$) model of the form

$y(t) = a + b(t+c)^{1/2}$, where slope = $y'(t) = \frac{1}{2} \cdot b(t+c)^{-1/2}$.

References for Frequency Drift

1. J.A. Barnes, "The Measurement of Linear Frequency Drift in Oscillators," *Proc. 15th Annu. PTTI Meeting*, pp. 551-582, Dec. 1983.
2. J.A. Barnes, "The Analysis of Frequency and Time Data," Austron, Inc., Dec. 1991.
3. M.A. Weiss and C. Hackman, "Confidence on the Three-Point Estimator of Frequency Drift," Proc. 24th Annu. PTTI Meeting, pp. 451-460, Dec. 1992.

5.23. All Tau

Stability calculations made at all possible tau values can provide an excellent indication of the variations in the results, and are a simple form of spectral analysis. In particular, cyclic variations are often the result of interference between the sampling rate and some periodic instability (such as environmental sensitivity). However, an all tau analysis is computationally intensive and can therefore be slow. For most purposes, however, it is not necessary to calculate values at every tau, but instead to do so at enough points to provide a nearly continuous curve on the display device (screen or paper). Such a "many tau" analysis can be orders of magnitude faster and yet provide the same information.
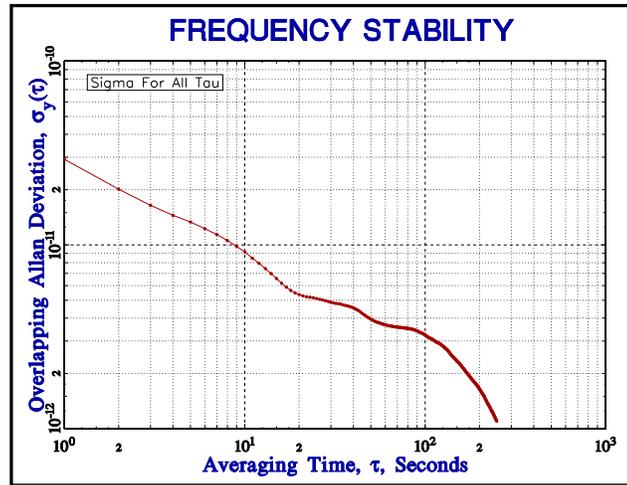
(a) All Tau Stability Plot
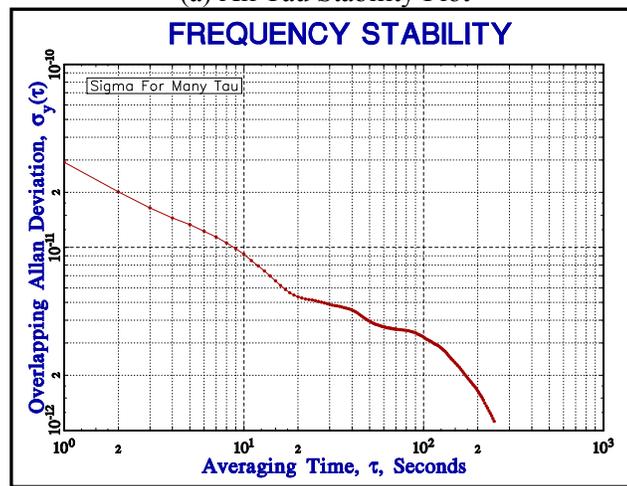
(b) Many Tau Stability Plot

Figure 28. Comparison of all tau and many tau stability.

5.25. Environmental Sensitivity

Environmental sensitivity should be treated separately from noise when one performs a stability analysis. However, it can be very difficult to distinguish between those different mechanisms for phase or frequency variations. It is often possible to control the environmental conditions sufficiently well during a stability run so that environmental effects such as temperature and supply voltage are negligible. Determining how well those factors have to be controlled requires knowledge of the device's environmental sensitivities. Those factors should be measured individually, if possible, over the largest reasonable excursions to minimize the effect of noise. Environmental sensitivity can best be determined by considering the physical mechanisms that apply within the unit under test. Useful information about the environmental sensitivity of frequency sources can be found in the references below. Some environmental factors affect phase and frequency differently, which can cause confusion. For example, temperature affects the phase delay through a cable. Dynamically, however, a temperature ramp produces a rate of change of phase that mimics a frequency change in the source. Because environmental sensitivity is highly dependent on device and application, it does not receive detailed consideration in this handbook. More information will be found in the following references.

References for Environmental Sensitivity

1. IEEE Standard 1193, *IEEE Guide for Measurement of Environmental Sensitivities of Standard Frequency Generators*.
2. MIL-0-55310, *General Specification for Military Specification, Oscillators, Crystal, Military Specifications and Standards*.
3. MIL-STD-810, *Environmental Test Methods and Engineering Guidelines*.
4. MIL-STD-202, *Test Methods for Electronic and Electrical Component Parts*.
5. C. Audion, N. Dimarcq, V. Giodano, J. Viennet, "Physical Origin of the Frequency Shifts in Cesium Beam Frequency Standards: Related Environmental Sensitivity," *Proceedings of the 22nd Annual Precise Time and Time Interval (PTTI) Applications and Planning Meeting*, pp. 419-440, 1990.
6. L.A. Breakiron, L.S. Cutler, H.W. Hellwig, J.R. Vig, G.M.R Winkler, "General Considerations in the Metrology of the Environmental Sensitivities of Standard Frequency Generators," *Proceedings of the 1992 IEEE Frequency Control Symposium*, pp. 816-830, 1992.
7. H. Hellwig, "Environmental Sensitivities of Precision Frequency Sources," *IEEE Transactions* IM-39, pp. 301-306, 1990.
8. E.M. Mattison, "Physics of Systematic Frequency Variations in Hydrogen Masers," *Proceedings of the 22nd Annual Precise Time and Time Interval (PTTI) Applications and Planning Meeting*, pp. 453-464, 1990.
9. W.J. Riley, "The Physics of Environmental Sensitivity of Rubidium Gas Cell Atomic Frequency Standards," *Proceedings of the 22nd Annual Precise Time and Time Interval (PTTI) Applications and Planning Meeting*, pp. 441-452, 1990.
10. F.L. Walls and J.J. Gagnepain, "Environmental Sensitivities of Quartz Crystal Oscillators," *IEEE Transactions UFFC-39*, no. 2, pp. 241-249, 1992.
11. P.S. Jorgensen, "Special Relativity and Intersatellite Tracking," *Navigation*, Vol. 35, No. 4, pp. 429-442, Winter 1988-89.
12. T.E. Parker, "Environmental Factors and Hydrogen Maser Frequency Sensitivity," *IEEE Trans. UFFC-46*, No. 3, pp.745-751, 1999.

5.26. Parsimony

In any measurement or analysis, it is desirable to minimize the number of extraneous parameters. This is not just a matter of elegance; the additional parameters may be implicit or arbitrary and thereby cause confusion or error if they are ignored or misunderstood. For example, the Allan deviation has the important advantage that, because of its convergence characteristics for all common clock noises, its expected value is independent of the number of data points. Many of the techniques used in the analysis of frequency stability, however, do require that certain parameters be chosen before they can be performed. For example, drift removal requires the choice of a model to which the data will be fit (perhaps using the criterion of white residuals). Outlier removal is an especially difficult case, where judgment often enters into the decision as to whether or not a datum is anomalous. A listing of some of these nonparsimonious parameters is given in Table 14:

Table 14. Non-Parsimonious Parameters in Frequency Stability Analysis

Type	Process	Parameter	Criterion	Remarks
Pre-processing	Outlier removal	Number of sigmas	Apply best judgment	Use of MAD-based robust statistics is recommended.
	Drift removal	Remove or not	Is drift deterministic? Is its cause known?	It is generally wise to remove deterministic drift before noise analysis
		Model	White residuals are sign that model is appropriate.	Model may have physical basis (e.g., diffusion process)
		Convergence limit	For iterative fit.	Generally uncritical – can hide deeply in algorithm
	Remove average frequency	Model	Noise type	Not necessarily simple arithmetic mean of frequency.
	Phase to frequency conversion	Tau	Accuracy	Is measurement interval known exactly? Is it the same for each point?
	Frequency to phase conversion	Tau	See above	
		Initial phase	Set to zero.	Generally arbitrary. Doesn't affect subsequent stability analysis.
		Normalize frequency	Use to emphasize noise rather than frequency offset	See "Remove average frequency above"
Analysis	Drift estimation	Model	Smallest residuals	Can be critical, especially for predictions. Known physics can help choose.
		Convergence limit	For iterative fit	Generally uncritical – can hide deeply in algorithm
	Frequency estimation	Noise model	Lowest residuals for known noise type	Generally uncritical
	Gaps	Skip, fill, omit, or exclude	Number, distribu-tion, noise type	Process choice affects results

		Allan deviation	Number of data points	Time available for measurement. Must remove outliers.	As many as possible
			Dead time	Property of measuring system	Avoid
			Maximum AF	Confidence degrades	As large as confidence allows
			Noise ID method to support error bars	Noise type significant affect on results	Generally unambiguous
		Hadamard deviation	All Allan deviation parameters	See above	
			To use rather than Allan deviation or separate drift removal	Easier handling of clock with drift or divergent noise	Commonly used in GPS control operations.
		Total deviation Thêo1	All Allan deviation parameters	See above	
			To use rather than Allan deviation	Better confidence at long tau.	Less commonly used
			Noise ID method to support bias removal	Critical for these biased estimators	Generally unambiguous
		Dynamic stability	Window and step size	Resolution, number of windows	Affects calc time
			Variance type	Data properties: noise, drift	AVAR or HVAR
			Viewpoint, mesh, color	Visibility	Personal preference
		Spectral analysis	Type – Parametric or non-parametric,	Convention	Analysis tools available Knowledge of analyst
			Windowing – Bias reduction	Clarity	Uncritical
			Smoothing – variance reduction	Clarity	Tradeoff vs resolution
			Presentation – plot or fit	Insight	Use both

Presenta-tion	Form	Table vs. Plot	Clarity	Use both
	Domain	Time or frequency	Clarity, correspondence with requirements	Use best
	Error bars	Include or not	Clarity	Error bars recommended
	Reference noise subtraction	Remove or not	If affects results	Only if similar to unknown
	Nominal vs. maximum at some confidence level	Cost/risk tradeoff	May be specified	Usually nominal at 1-sigma confidence
	Notation	State outlier & drift removal, environmental conditions, etc.	Judgment of analyst	Disclose choices

5.27. Transfer Functions

Variances can be related to the spectral density of the fractional frequency fluctuations by their transfer functions. For example, for the Hadamard variance, this relationship is

$$\sigma^2{}_H(\tau) = \int_0^{f_h} S_y(f) \cdot \left| H_H(f) \right|^2 df , \qquad (57)$$

where $\sigma^2{}_H(\tau)$ is the three-sample, zero dead-time, binominally weighted Hadamard variance, $S_y(f)$ is the spectral density of the fractional frequency fluctuations, $H_H(f)$ is its transfer function, and f_h is the upper cutoff frequency (determined by hardware factors). The transfer function is determined by the statistic's time-domain sampling pattern.

The transfer functions for the most common variances used in frequency stability analysis are shown in Table 15:

Table 15. Transfer functions for the most common variances.

Variance	Magnitude Squared Transfer Function $\|H(f)\|^2$
Allan	$2\left(\dfrac{\sin \pi\tau f}{\pi\tau f}\right)^2 \sin^2 \pi\tau f$
Hadamard	$2^4\left(\dfrac{\sin \pi\tau f}{\pi\tau f}\right)^2 \sin^4 \pi\tau f$

For $\pi\tau f \ll 1$, the transfer function of the Allan variance behaves as $(\pi\tau f)^2$, indicating that it is convergent for power law processes S_y^α down to as low as $\alpha = -2$ (Random Walk FM), while the transfer function of the Hadamard variance behaves as $(\pi\tau f)^4$, indicating that it is convergent for power law processes as low as $\alpha = -4$ (Random Run FM).

The squared magnitudes of these transfer functions are shown in the plots below:
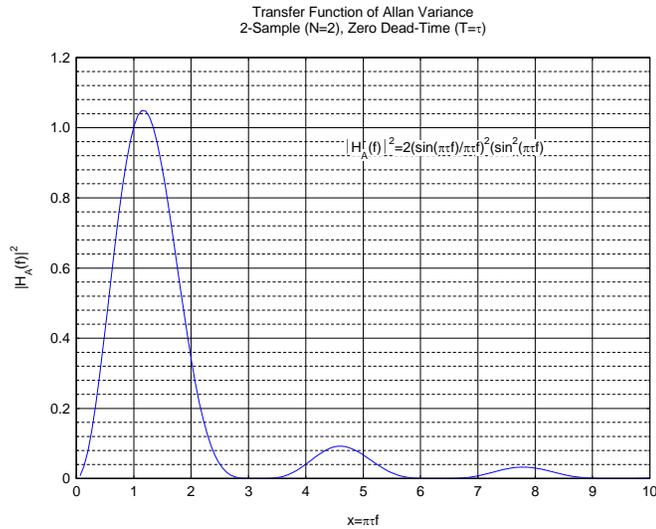
(a) Allan Variance
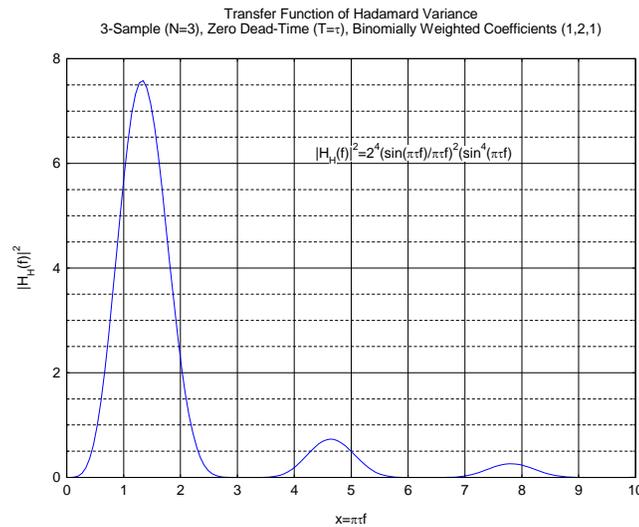
(b) Hadamard Variance

Figure 29. Squared magnitudes of transfer functions for (a) Allan Variance and (b) Hadamard Variance.

These responses have their peaks where the frequency is one-half the sampling rate, and nulls where it is a multiple of the sampling rate (i.e., at $f = n/\tau$, where n is an integer). As a spectral estimator, the Hadamard variance has slightly higher resolution than the Allan variance, since the equivalent noise bandwidths of the Hadamard and Allan spectral windows are $0.411\ \tau^{-1}$ and $0.476\ \tau^{-1}$, respectively [5].

Similar transfer functions exist for the modified, total, and Thêo1 variances.

References for Transfer Functions

1. J. Rutman, "Oscillator Specifications: A Review of Classical and New Ideas," *1977 IEEE International Freq. Contrl. Symp.*, pp. 291-301, June 1977.
2. J. Rutman, "Characterization of Phase and Frequency Instabilities in Precision Frequency Sources," *Proc. IEEE, Vol. 66, No. 9*, September1978.
3. R.G. Wiley, "A Direct Time-Domain Measure of Frequency Stability: The Modified Allan Variance," *IEEE Trans. Instrum. Meas.*, **IM-26**, pp. 38-41.
4. D.A. Howe and C.A. Greenhall, "Total Variance: A Progress Report on a New Frequency Stability Characterization," *Proc. 29th Annual Precise Time and Time Interval (PTTI) Meeting*, December 1977, pp. 39-48.
5. D.A. Howe, "ThêoH: A Hybrid, High-Confidence Statistic that Improves on the Allan Deviation," *Metrologia* **43** (2006), S322-S331.
6. K. Wan, E. Visr and J. Roberts, "Extended Variances and Autoregressive Moving Average Algorithm for the Measurement and Synthesis of Oscillator Phase Noise," *Proc. 43rd Annu. Freq. Contrl. Symp.*, June 1989, pp. 331-335.

6   Frequency Domain Stability

Frequency stability can also be characterized in the frequency domain in terms of a power spectral density (PSD) that describes the intensity of the phase or frequency fluctuations as a function of Fourier frequency. Spectral stability measures are directly related to the underlying noise processes, and are particularly appropriate when the phase noise of the source is of interest.

> Frequency domain stability measures are based on power spectral densities that characterize the intensity of the phase or frequency fluctuations as a function of Fourier frequency.

6.1.   Noise Spectra

The random phase and frequency fluctuations of a frequency source can be modeled by power law spectral densities of the form

$$S_y(f) = h(\alpha)f^{\alpha},$$

where: $S_y(f)$ =   one-sided power spectral density of y, the
                            fractional frequency fluctuations, 1/Hz
        f   =   Fourier or sideband frequency, hertz,
       $h(\alpha)$ =   intensity coefficient, and
        α =   exponent of the power law noise process.

The most commonly encountered noise spectra are

White (f^{0})
Flicker (f^{-1})
Random Walk (f^{-2})
Flicker Walk (f^{-3}) .

Examples of these noise types are shown in the figure below.
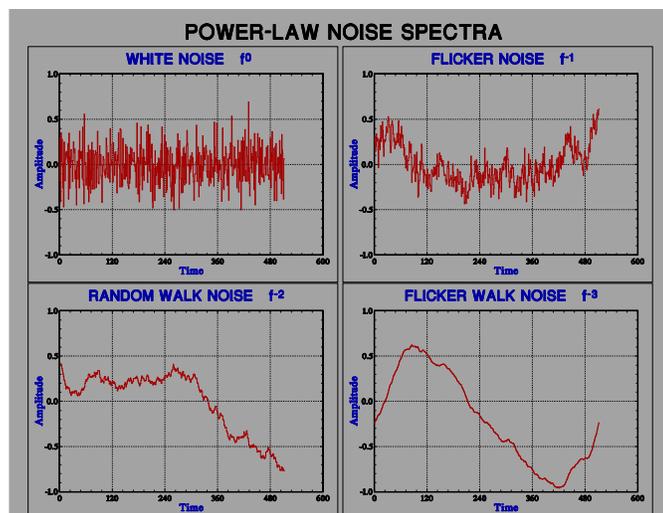
Figure 30. Examples of noise types.

Power law spectral models can be applied to both phase and frequency power spectral densities. Phase is the time integral of frequency, so the relationship between them varies as $1/f^2$:

$$S_x(f) = \frac{S_y(f)}{(2\pi f)^2},$$ (59)

where $S_x(f)$ = PSD of the time fluctuations, sec²/Hz. (60)

Two other quantities are also commonly used to measure phase noise:

$S_\phi(f)$ = PSD of the phase fluctuations, rad²/Hz and its logarithmic equivalent £(f), dBc/Hz.

The relationship between these is

$$S_\phi(f) = (2\pi\nu_0)^2 \cdot S_x(f) = \left(\frac{\nu_0}{f}\right)^2 \cdot S_y(f),$$ (61)

and

$$\mathcal{L}(f) = 10 \cdot \log\left[\frac{1}{2} \cdot S_y(f)\right],$$ (62)

where ν_0 is the carrier frequency, hertz.

The power law exponent of the phase noise power spectral densities is $\beta = \alpha - 2$. These frequency-domain power law exponents are also related to the slopes of the following time domain stability measures:

Allan variance $\sigma^2_y(\tau)$ $\mu = -(\alpha+1), \alpha<2$
Modified Allan variance   Mod $\sigma^2_y(\tau)$ $\mu' = -(\alpha+1, \alpha<3$
Time variance $\sigma^2_x(\tau)$ $\eta = -(\alpha-1), \alpha<3$.

The spectral characteristics of the power law noise processes commonly used to describe the performance of frequency sources are shown in the following table:

Table 16. Spectral characteristics of power law noise processes

Noise Type	α	β	μ	μ'	η
White PM	2	0	–2	–3	–1
Flicker PM	1	–1	–2	–2	0
White FM	0	–2	–1	–1	1
Flicker FM	–1	–3	0	0	2
Random walk FM	–2	–4	1	1	3

6.2. Power Spectral Densities

Four types of power spectral density are commonly used to describe the stability of a frequency source:

PSD of Frequency Fluctuations $S_y(f)$
The power spectral density of the fractional frequency fluctuations y(t) in units of 1/Hz is given by $S_y(f) = h(\alpha) \cdot f^\alpha$, where f = sideband frequency, Hz.

PSD of Phase Fluctuations $S_\phi(f)$

68

The power spectral density of the phase fluctuations in units of rad²/Hz is given by $S_\phi(f) = (2\pi\nu_0)^2 \cdot S_x(f)$, where $\nu_0 =$ carrier frequency, Hz.

PSD of Time Fluctuations $S_x(f)$

The power spectral density of the time fluctuations x(t) in units of sec²/Hz is given by $S_x(f) = h(\beta) \cdot f^\beta = S_y(f)/(2\pi f)^2$, where: $\beta = \alpha–2$. The time fluctuations are related to the phase fluctuations by $x(t) = \phi(t)/2\pi\nu_0$. Both are commonly called "phase" to distinguish them from the independent time variable, t.

SSB Phase Noise £(f)

The SSB phase noise in units of dBc/Hz is given by $£(f) = 10 \cdot \log [\frac{1}{2} \cdot S_\phi(f)]$. This is the most common function used to specify phase noise.

6.3. Phase Noise Plot

The following diagram shows the slope of the SSB phase noise, £(f), dBc/Hz versus log f, Fourier frequency, Hz for various power law noise processes.
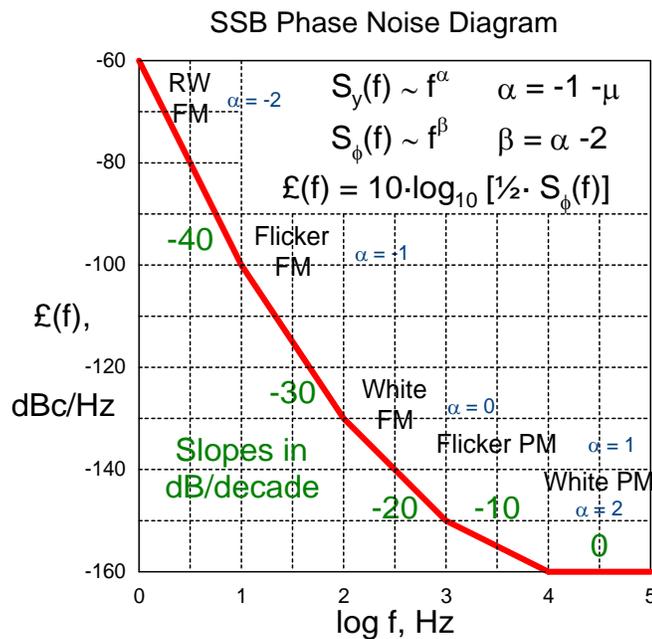
Figure 31. SSB Phase Noise Plot.

6.4. Spectral Analysis

Spectral analysis is the process of characterizing the properties of a signal in the frequency domain, either as a power spectral density for noise, or as the amplitude and phase at discrete frequencies. Spectral analysis can thus be applied to both noise and discrete components for frequency stability analysis. For the former, spectral analysis complements statistical analysis in the time domain. For the latter, spectral analysis can aid in the identification of periodic components such as interference and environmental sensitivity. Time domain data can be used to perform spectral analysis via the Fast Fourier Transform (FFT), and there is much technical literature on that subject [2,3]. While, in principle, time and frequency domain analyses provide equivalent information, in practice, one or the other is usually

preferred, either because of measurement and/or analysis convenience, or because the results are more applicable to a particular application. Spectral analysis is most often used to characterize the short-term (< 1 s) fluctuations of a frequency source as a plot of phase noise power spectral density versus sideband frequency, while a time domain analysis is most often used to provide information about the statistics of its instability over longer intervals (> 1 s). Modern instrumentation is tending to merge these approaches by digitizing the signal waveform at high sampling rates, thereby allowing FFT analysis at relatively high Fourier frequencies. Nevertheless, there are many pitfalls and subtleties associated with spectral analysis that must be considered if meaningful results are to be obtained.

6.5. PSD Windowing

Data windowing is the process of applying a weighting function that falls off smoothly at the beginning and end to avoid spectral leakage in an FFT analysis. Without windowing, bias will be introduced that can severely restrict the dynamic range of the PSD result. The most common windowing types are Hanning, Hamming, and Multitaper. The classic Hanning and Hamming windows can be applied more than one time.

6.6. PSD Averaging

Without filtering or averaging, the variances of the PSD results are always equal to their values regardless of the size of the time domain data set. More data provide finer frequency resolution, not lower noise (while the data sampling time determines the highest Fourier frequency). Without averaging, for white noise, each spectral result has only two degrees of freedom. Some sort of filtering or averaging is usually necessary to provide less noise in the PSD results. This can be accomplished by dividing the data into sections, performing an FFT analysis on each section separately, and then averaging those values to obtain the final PSD result. The averaging factor improves the PSD standard deviation by the square root of the averaging factor. The tradeoff in this averaging process is that each section of the data is shorter, yielding a result with coarser frequency resolution that does not extend to as low a Fourier frequency.

6.7. Multitaper PSD Analysis

The multitaper PSD analysis method offers a better compromise among bias, variance and spectral resolution. Averaging is accomplished by applying a set of orthogonal windowing (tapering) functions called discrete prolate spheroidal sequences (DPSS) or Slepian functions to the entire data array. An example of seven of these functions for order J=4 is shown in Figure 32.
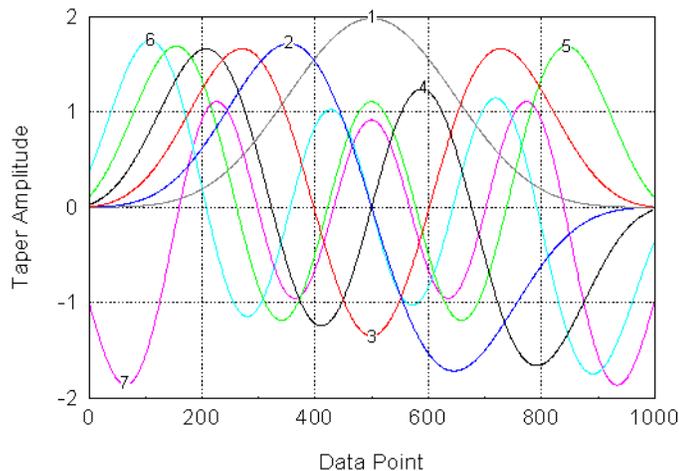
Figure 32. Slepian SPSS taper functions, J=4, #=7.

The first function resembles a classic window function, while the others sample other portions of the data. The higher windows have larger amplitude at the ends that compensates for the denser sampling at the center. These multiple tapering functions are defined by two parameters, the order of the function, J, which affects the resolution bandwidth, and the number of windows, which affects the variance. A higher J permits the use of more windows without introducing bias, which provides more averaging (lower variance) at the expense of lower spectral resolution, as shown Table 14:

Table 17. The two parameters of the tapering function.

Order J	# Windows
2.0	1-3
2.5	1-4
3.0	1-5
3.5	1-6
4.0	1-7
4.5	1-8
5.0	1-9

The resolution BW is given by 2J/Nt, where N is the number of data points sampled at time interval t. An adaptive algorithm can be used to weight the contributions of the individual tapers for lowest bias. The multitaper PSD has a flat-topped response for discrete spectral components that is nevertheless narrower than an averaged periodogram with the same variance. It is therefore particularly useful for examining discrete components along with noise.

6.8. PSD Notes

A carrier frequency parameter applies to the $S_\phi(f)$ and £(f) PSD types. The number of Fourier frequency points is always the power of 2 greater than or equal to one-half of the number of time domain data points, n. The spacing between Fourier frequency points is 1/nt, and the highest Fourier frequency is 1/2t. If averaging is done, the value of n is reduced by the averaging factor. The PSD fit is a least-squares power law line through octave-band PSD averages [6].

For characterizing frequency stability, a spectral analysis is used primarily for the analysis of noise (not discrete components), and should include the quantitative display of power law noise in common PSD units, perhaps with fits to integer power law noise processes. Amplitude corrections need to be made for the noise response of the windowing functions. The amplitude of discrete components should be increased by the log of the BW (Fourier frequency spacing in hertz), which is a negative number for typical sub-hertz bandwidths.

References for Spectral Analysis

1. W.H. Press, B.P. Flannery, S.A. Teukolsky and W.T. Vetterling, <u>Numerical Recipes in C</u>, Cambridge University Press, 1988, ISBN 0-521-35465-X, Chapter 12.
2. D.J. Thomson, "Spectrum Estimation and Harmonic Analysis," *Proc. IEEE* , Vol. 70, No. 9, Sept. 1982, pp. 1055-1096.
3. D.B. Percival and A.T. Walden, **Spectral Analysis for Physical Applications**, Cambridge University Press, 1993, ISBN 0-521-43541-2.
4. J. M. Lees and J. Park, "A C-Subroutine for Computing Multi-Taper Spectral Analysis," *Computers in Geosciences* , Vol. 21, 1995, pp. 195-236.
5. Help File for the *AutoSignal* program for spectral analysis, AISN Software, Inc., 1999.
6. J. McGee and D.A. Howe, "ThêoH and Allan Deviation as Power-Law Noise Estimators," *IEEE Trans. Ultrasonics, Ferroelectrics and Freq. Contrl.*, Feb. 2007.

7   Domain Conversions

The stability of a frequency source can be specified and measured in either the time or the frequency domain. Examples of these stability measures are the Allan variance, $\sigma^2_y(\tau)$, in the time domain, and the spectral density of the fractional frequency fluctuations, $S_y(f)$, in the frequency domain. Conversions between these domains may be made by numerical integration of their fundamental relationship, or by an approximation method based on a power law spectral model for the noise processes involved. The latter method can be applied only when the dominant noise process decreases toward higher sideband frequencies. Otherwise, the more fundamental method based on numerical integration must be used. The general conversion from time to frequency domain is not unique because white and flicker phase noise have the same Allan variance dependence on τ. When performing any of these conversions, it is necessary to choose a reasonable range for σ and τ in the domain being converted to. The main lobe of the $\sigma_y(\tau)$ and Mod $\sigma_y(\tau)$ responses occur at the Fourier frequency $f = 1/2\tau$.

> The stability of a frequency source can be specified and measured in either the time or the frequency domain. One domain is often preferred to specify the stability because it is most closely related to the particular application. Similarly, stability measurements may be easier to perform in one domain than the other. Conversions are possible between these generally equivalent measures of frequency stability.

Time domain frequency stability is related to the spectral density of the fractional frequency fluctuations by the relationship

$$\sigma^2(M,T,\tau) = \int_0^\infty S_y(f) \cdot |H(f)|^2 \cdot df \,, \tag{63}$$

where $|H(f)|^2$ is the transfer function of the time domain sampling function.

The transfer function of the Allan (two-sample) time domain stability is given by

$$|H(f)|^2 = 2\left[\frac{\sin^4(\pi\tau f)}{(\pi\tau f)^2}\right], \tag{64}$$

and therefore the Allan variance can be found from the frequency domain by the expression

$$\sigma^2_y(\tau) = 2\int_0^{f_h} S_y(f)\frac{\sin^4(\pi\tau f)}{(\pi\tau f)^2}df \,, \tag{65}$$

The equivalent expression for the modified Allan variance is

$$Mod\,\sigma^2_y(\tau) = \frac{2}{N^4\pi^2\tau_0^2}\int_0^{f_h}\frac{S_y(f)\sin^6(\pi\tau f)}{f^2\sin^2(\pi\tau_0 f)}df \,. \tag{66}$$

7.1. Power Law Domain Conversions

Domain conversions may be made for power law noise models by using the following equation and conversion formulae:

$$\sigma^2_y(\tau) = h_{-2}\frac{(2\pi)^2}{6}\tau + h_{-1}2\log_e 2 + h_0\frac{1}{2\tau} + h_1\frac{1.038+3\log_e 2\pi f_h\tau}{2\pi^2\tau^2} + h_2\frac{3f_h}{(2\pi)^2\tau^2} \,, \tag{67}$$

where the h_α terms define the level of the various power law noises.

Noise Type	$\sigma^2_y(\tau)$	$S_y(f)$
RW FM	$A \cdot f^{2} \cdot S_y(f) \cdot \tau^{1}$	$A^{-1} \cdot \tau^{-1} \cdot \sigma^2_y(\tau) \cdot f^{-2}$
F FM	$B \cdot f^{1} \cdot S_y(f) \cdot \tau^{0}$	$B^{-1} \cdot \tau^{0} \cdot \sigma^2_y(\tau) \cdot f^{-1}$
W FM	$C \cdot f^{0} \cdot S_y(f) \cdot \tau^{-1}$	$C^{-1} \cdot \tau^{1} \cdot \sigma^2_y(\tau) \cdot f^{0}$
F PM	$D \cdot f^{-1} \cdot S_y(f) \cdot \tau^{-2}$	$D^{-1} \cdot \tau^{2} \cdot \sigma^2_y(\tau) \cdot f^{1}$
W PM	$E \cdot f^{-2} \cdot S_y(f) \cdot \tau^{-2}$	$E^{-1} \cdot \tau^{2} \cdot \sigma^2_y(\tau) \cdot f^{2}$

where
$A = 4\pi^2/6$
$B = 2 \cdot \ln 2$
$C = 1/2$
$D = 1.038 + 3 \cdot \ln(2\pi f_h \tau_0)/4\pi^2$
$E = 3 f_h /4\pi^2$

and f_h is the upper cutoff frequency of the measuring system in hertz, and τ_0 is the basic measurement time. This factor applies only to white and flicker PM noise. The above conversion formulae apply to the ThêoH hybrid statistic as well as to the Allan variance.

7.2. Example of Domain Conversions

This section shows an example of time and frequency domain conversions. First, a set of simulated power law noise data is generated, and the time domain properties of this noise are analyzed by use of the overlapping Allan deviation. Next, the same data are analyzed in the frequency domain with an £(f) PSD. Then, a power law domain conversion is done, and those results are compared with those of the spectral analysis. Finally, the other power spectral density types are examined.

For this example, we generate 4097 phase data points of simulated white FM noise with a one-second Allan deviation value $\sigma_y(1) = 1 \times 10^{-11}$ and a sampling interval $\tau = 1$ ms. The number of points is chosen as an even power of 2 for the frequency data. The generated set of simulated white FM noise is shown as frequency data in Figure 33a, and their overlapping Allan deviation is shown in Figure 33b. The $\sigma_y(1)$ white FM noise fit parameter is 1.08e–11, close to the desired value.
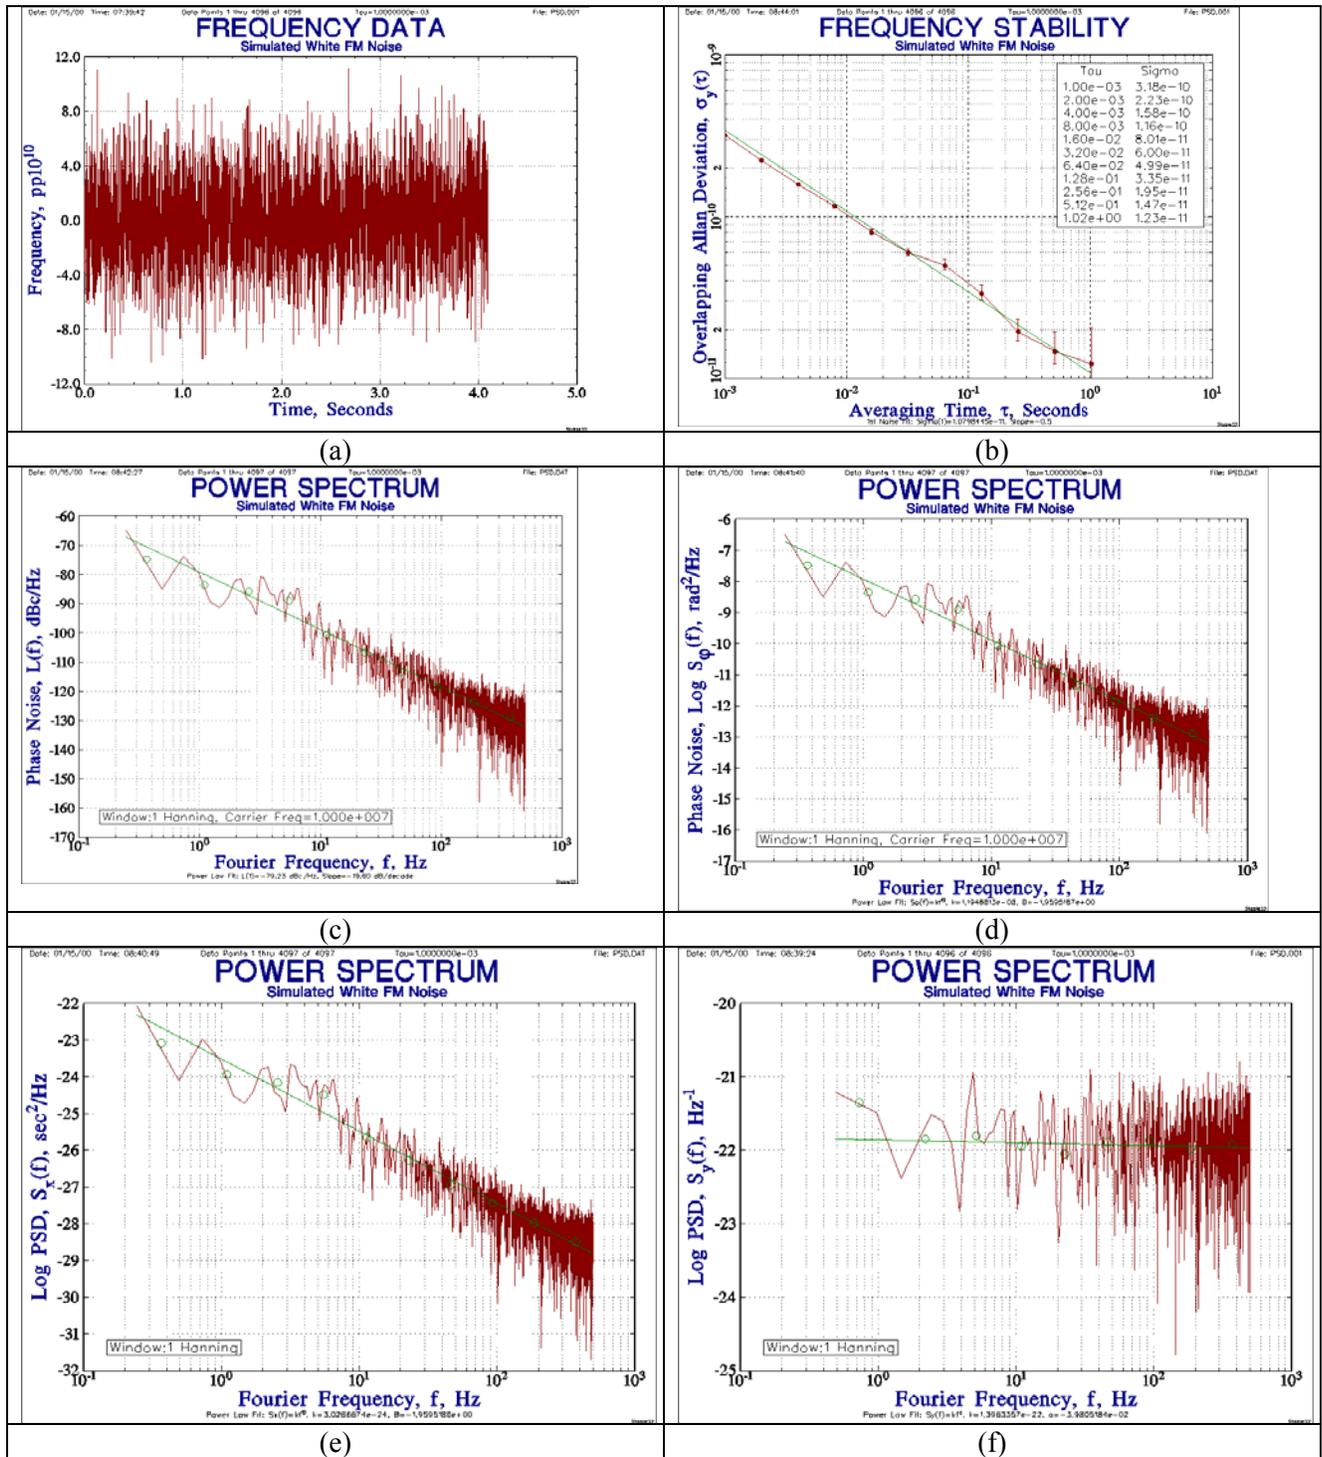
Figure 33. (a) Simulated W FM Noise Data. (b) Overlapping Allan Deviation. (c) £(f) Power Spectral Density, (d) $S_\phi(f)$ Power Spectral Density. (e) $S_x(f)$ Power Spectral Density. (f) $S_y(f)$ Power Spectral Density

The power spectrum for the phase data is calculated by use of a 10 MHz carrier frequency and a £(f) power spectral density type, the SSB phase noise to signal ratio in a 1 Hz BW as a function of sideband frequency, f, as shown in Figure 33c. The fit parameters show an £(1) value of –79.2 dBc/Hz and a slope of –19.6 dB/decade, in close agreement with the expected values of –80 dBc/Hz and –20 dB/decade.

The expected PSD values that correspond to the time domain noise parameters used to generate the simulated power-law noise can be determined by the power law domain conversion formulae of Section 7.1, as shown in Table18.

Table 18. Domain calculation results

```
Frequency Domain:        PSD Type: L(f), dBc/Hz
                         SB Freq (Hz): 1.00000e+00
                         Carrier(MHz): 1.00000e+01

Time Domain:             Sigma Type: Normal
                         Tau (Sec): 1.00000e-03
                         Avg Factor: 1000

Power Law Noise:

Type   dB/dec PSD     Type   Mu   Sigma
RWFM   -40    None    RWFM   +1   0.00000e+00
FFM    -30    None    FFM     0   0.00000e+00
WFM    -20    -80.0   WFM    -1   1.00000e-11
FPM    -10    None    FPM    -2   0.00000e+00
WPM     0     None    WPM    -2   0.00000e+00
All           -80.0   All         1.00000e-11
```

The other types of PSD that are commonly used for the analysis of frequency domain stability analysis are $S_\phi(f)$, the spectral density of the phase fluctuations in rad^2/Hz; $S_x(f)$, the spectral density of the time fluctuations in sec^2/Hz; and $S_y(f)$, the spectral density of the fractional frequency fluctuations in units of 1/Hz. The expected value of all these quantities for the simulated white FM noise parameters with $\sigma_y(1) = 1.00e{-}11$, $\tau = 1.00e{-}3$, and $f_o = 10$ MHz are shown in the following table.

Table 19. Expected value of PSD types for the simulated white FM noise parameters.

Parameter	PSD Type			
	£(f), dBc/Hz	$S_\phi(f)$, rad^2/Hz	$S_x(f)$, sec^2/Hz	$S_y(f)$, 1/Hz
Data Type	Phase	Phase	Phase	Frequency
Simulated Value	−80	2e–8	5.066e–24	2e–22
Log Value	same	−7.70	−23.3	−21.7
Slope, dB/decade	−20	−20	−20	0
Fit Value	−79.2	1.19e–8	3.03e–24	1.40e–22
Fit Exponent	−1.96	−1.96	−1.96	−0.04

References for Domain Conversion

1. A.R. Chi, "The Mechanics of Translation of Frequency Stability Measures Between Frequency and Time Domain Measurements," *Proc. 9th PTTI Meeting*, pp. 523-548, Dec.1977.
2. J. Rutman, "Relations Between Spectral Purity and Frequency Stability," *Proc. 28th Freq. Contrl. Symp.*, pp. 160-165., May 1974.
3. R. Burgoon and M.C. Fisher, "Conversion Between Time and Frequency Domain of Intersection Points of Slopes of Various Noise Processes," *32th Freq. Contrl. Symp.*, pp.514-519, June 1978.
4. W.F. Egan, "An Efficient Algorithm to Compute Allan Variance from Spectral Density," *IEEE Trans. Instrum. Meas.*, Vol. 37, No. 2, pp. 240-244, June 1988.
5. F. Vernotte, J. Groslambert and J.J. Gagnepain, "Practical Calculation Methods of Spectral Density of Instantaneous Normalized Frequency Deviation from Instantaneous Time Error Samples," *Proc. 5th European Freq. and Time Forum*, pp. 449-455, March 1991.
6. C.A. Greenhall, "Does Allan Variance Determine the Spectrum?" *Proc. 1997 Intl. Freq. Cont. Symp.*, pp. 358-365, June 1997.
7. F. Thomson, S. Asmar and K. Oudrhiri, "Limitations on the Use of the Power-Law Form of Sy(f) to Compute Allan Variance," *IEEE Trans. UFFC*, Vol. 52, No. 9, pp. 1468-1472, Sept. 2005.

8   Noise Simulation

It is valuable to have a means of generating simulated power law clock noise having the desired noise type (white phase, flicker phase, white frequency, flicker frequency, and random walk frequency noise), Allan deviation, frequency offset, frequency drift, and perhaps a sinusoidal component. This can serve as both a simulation tool and as a way to validate stability analysis software, particularly for checking numerical precision, noise recognition, and modeling. A good method for power-law noise generation is described in Reference 8. The noise type and time series of a set of simulated phase data are shown in Table 20:

Table 20. Noise type and time series for a set of simulated phase data.

Noise Type	Phase Data Plot
Random walk FM $\alpha = -2$ Random run noise	
Flicker FM $\alpha = -1$ Flicker walk noise	
White FM $\alpha = 0$ Random walk noise	
Flicker PM $\alpha = 1$ Flicker noise	
White PM $\alpha = 2$ White noise	
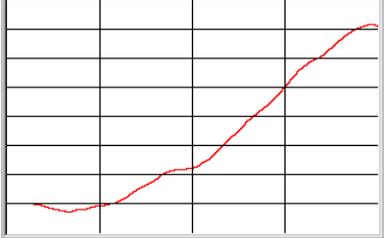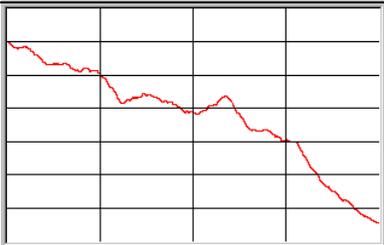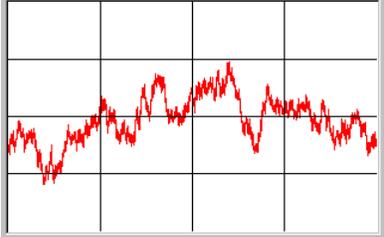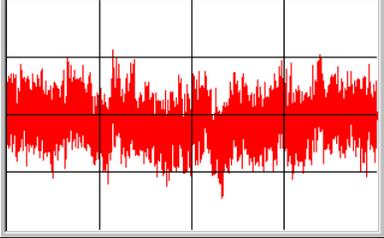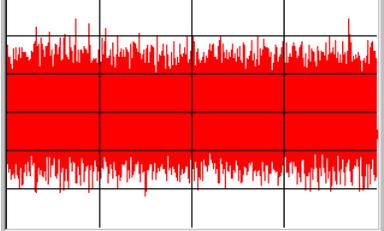
8.1. White Noise Generation

White noise generation is straightforward. One popular technique is to first generate two independent uniformly distributed random sequences [1], and combine them using the Box-Muller transform [2,3] to produce a white spectrum with Gaussian deviates. Another method is to generate 12 independent random sequences uniformly distributed between 0 and 1, add them, and subtract 6 [4]. This will, via the central limit theorem, produce a Gaussian distribution having zero mean and unit variance. White noise can be numerically integrated and differenced to transform it by $1/f^2$ and f^2, respectively, to produce simulated noise having any even power law exponent.

8.2. Flicker Noise Generation

Flicker noise is more difficult to generate because it cannot be described exactly by a rational transfer function, and much effort has been devoted to generating it [5-9]. The most common methods involve linear filtering by RC ladder networks [5], or by FFT transformation [7,9]. The FFT method can produce noise having any integer power law exponent from $\alpha = -2$ (RW FM) to $\alpha = +2$ (W PM) [7, 8].

8.3. Flicker Walk and Random Run Noise Generation

The more divergent flicker walk FM (a = –3) and random run FM (a = –4) power law noise types may be generated by using the $1/f^2$ spectral property of a frequency to phase conversion. For example, to generate RR FM noise, first generate a set of RW FM phase data. Then treat this RW FM phase data as frequency data, and convert it to a new set of RR FM phase data.

8.4. Frequency Offset, Drift, and Sinusoidal Components

Beside the generation of the desired power law noise, it is desirable to include selectable amounts of frequency offset, frequency drift, and a sinusoidal component in the simulated clock data.

References for Noise Simulation

1. S.K. Park and K.W. Miller, "Random Number Generators: Good Ones are Hard to Find," *Comm. ACM*, Vol. 31, No. 10, pp. 1192-1201.
2. G.E.P Box and M.E. Muller, "A Note on the Generation of Random Normal Deviates," *Annals Math. Stat*, V. 29, pp. 610-611,1958.
3. W.H. Press, B.P. Flannery, S.A. Teukolsky and W.T. Vetterling, Numerical Recipes in C, Cambridge Univ. Press, Cambridge, U.K., 1988.
4. J.A. Barnes, "The Analysis of Frequency and Time Data," Austron, Inc., Dec. 1991.
5. J.A. Barnes and D.W. Allan, "A Statistical Model of Flicker Noise," *Proceedings of the IEEE*, Vol. 54, No. 2, pp. 176-178, February 1966.
6. J.A. Barnes, "The Generation and Recognition of Flicker Noise," *NBS Report 9284*, U.S Department of Commerce, National Bureau of Standards, June 1967.
7. C.A. Greenhall and J.A. Barnes, "Large Sample Simulation of Flicker Noise," *Proc. 19th Annu. PTTI Meeting*, pp. 203-217, Dec. 1987 and *Proc. 24th Annu. PTTI Meeting*, p. 461, Dec. 1992.
8. N.J. Kasdin and T. Walter, "Discrete Simulation of Power Law Noise," *Proc. 1992 IEEE Freq. Contrl. Symp.*, pp. 274-283, May 1992.
9. S. Bregni, L. Valtriani, F. Veghini, "Simulation of Clock Noise and AU-4 Pointer Action in SDH Equipment," *Proc. of IEEE GLOBECOM '95*, pp. 56-60, Nov. 1995.
10. S. Bregni, "Generation of Pseudo-Random Power-Law Noise Sequences by Spectral Shaping," *Communications World*, Editor: N. Mastorakis, WSES Press, 2001.
11. C.A. Greenhall, "FFT-Based Methods for Simulating Flicker FM," *Proceedings of the 34th Annual Precise Time and Time Interval (PTTI) Systems and Applications Meeting*, 2002.

9   Measuring Systems

Frequency measuring systems are instruments that accept two or more inputs, one of which may be considered to be the reference, and compare their relative phase or frequencies. These systems can take many forms, from the direct use of a frequency counter to elaborate low-noise, high-resolution multichannel clock measuring systems with associated archival databases. They can be custom built or bought from several organizations specializing in such systems. The most

> A frequency measuring system with adequate resolution and a low noise floor is necessary to make precision clock measurements.

important attribute of a frequency measuring system is its resolution, which, for high performance devices, requires 1 ps/s (pp10^{12}) or better resolution, and more elaborate hardware than a counter. The resolution of a digital frequency, period, or time interval counter is determined mainly by its speed (clock rate) and the performance of its analog interpolator (if any). That resolution generally improves linearly with the averaging time of the measurement, and it can be enhanced by preceding it with a mixer that improves the resolution by the heterodyne factor, the ratio of the RF input to the IF beat frequencies. Noise is another important consideration for a high-performance measuring system whose useful resolution may be limited by its noise floor, the scatter in the data when the two inputs are driven coherently by the same source. The performance of the measuring system also depends on the stability of its reference source. A low noise ovenized quartz crystal oscillator may be the best choice for a reference in the short term (1 to 100 s), while a active hydrogen maser generally provides excellent stability at averaging times out to several days, and cesium beam tube devices at longer averaging times.

Three methods are commonly used for making precise time and frequency measurements, as described below.

9.1.   Time Interval Counter Method

The time interval counter method divides the two sources being compared down to a much lower frequency (typically 1 pulse/second) and measures their time difference with a high resolution time interval counter:
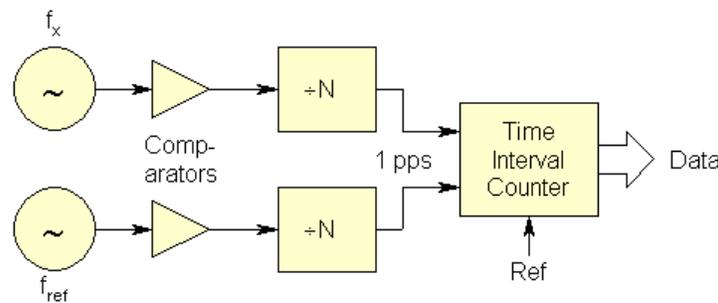
Figure 34. Block diagram of a time interval counter measuring system.

This measurement method is made practical by modern high-resolution interpolating time interval counters that offer 10 digit/s resolution. The resolution is not affected by the division ratio, which sets the minimum measurement time, and determines how long data can be taken before a phase spillover occurs (which can be hard to remove from a data set). A source having a frequency offset of 1×10^{-6} can, for example, be measured for only about 5.8 days before a 1 pps phase spillover occurs after being initially set at the center. Drift in the trigger point of the counter can be a limitation to this measurement method.

9.2.   Heterodyne Method

The heterodyne method mixes (subtracts) the two sources being compared, and measures the period of the resulting audio-frequency beat note. The measurement resolution is increased by the heterodyne factor (the ratio of the carrier to the beat frequency).
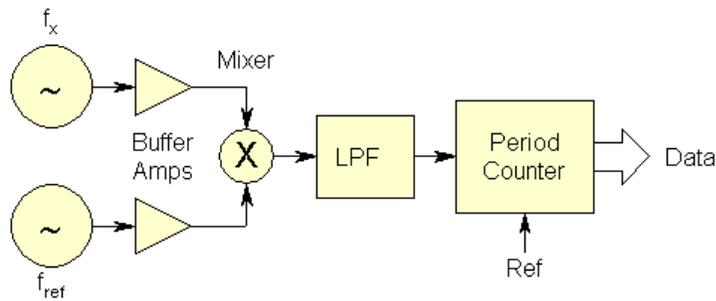
Figure 35. Block diagram of a heterodyne measuring system.

This heterodyne technique is a classic way to obtain high resolution with an ordinary period counter. It is based on the principle that phase information is preserved in a mixing process. For example, mixing a 10 MHz source against a 9.9999 MHz Hz offset reference will produce a 100 Hz beat signal whose period variations are enhanced by a factor of 10 MHz/100 Hz = 10^5. Thus a period counter with 100 ns resolution (10 MHz clock) can resolve clock phase changes of 1 ps. A disadvantage of this approach is that a stable offset reference is required at exactly the right frequency. Even worse, it can measure only frequency, requires a priori knowledge of the sense of the frequency difference, and often has dead time between measurements.

9.3. Dual Mixer Time Difference Method

The third method, in effect, combines the best features of the first two, using a time interval counter to measure the relative phase of the beat signals from a pair of mixers driven from a common offset reference:
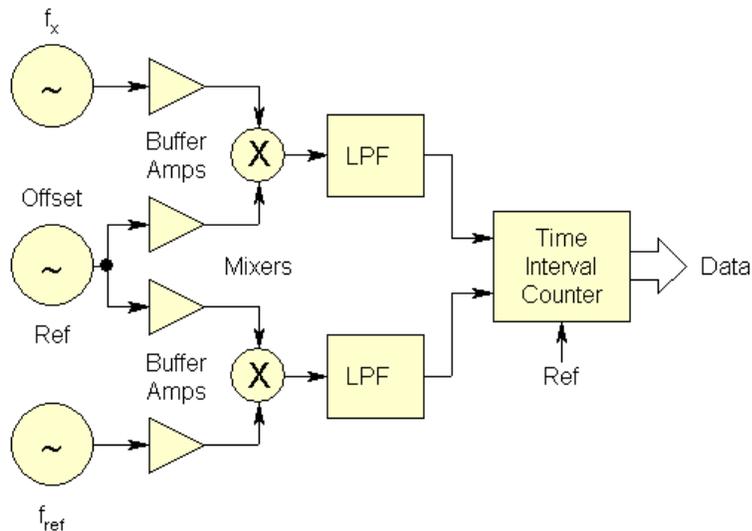
Figure 36. Block diagram of a dual mixer time difference measuring system.

This dual mixer time difference (DMTD) setup is arguably the most precise way of measuring an ensemble of clocks all having the same nominal frequency. When expanded to multiple channels by adding additional buffer amplifiers and mixers, and time tagging the zero-crossings of the beat notes for each channel, this arrangement allows any two of the clocks to be intercompared. The offset reference need not be coherent, nor must it have particularly low noise or high accuracy, because its effect cancels out in the overall measurement process. For best cancellation, the zero-crossings should be coincident or interpolated to a common epoch. Additional counters can be used to count the whole beat note cycles to eliminate their ambiguity, or the zero-crossings can simply be time tagged. The measuring system resolution is determined by the time interval counter or time-tagging hardware, and the mixer heterodyne

factor. For example, if two 5 MHz sources are mixed against a common 5 MHz to 10 Hz offset oscillator (providing a $5 \times 10^6/10 = 5 \times 10^5$ heterodyne factor), and the beat note is time tagged with a resolution of 100 ns(10 MHz clock), the measuring overall system resolution is $10^{-7}/5 \times 10^5 = 0.2$ ps.

Multichannel DMTD clock measuring systems have been utilized by leading national and commercial metrology laboratories for a number of years [1-5]. An early commercial version is described in Reference [3], and a newer technique is described in Reference [8]. A direct digital synthesizer (DDS) can be used as the offset reference to allow measurements to be made at any nominal frequency within its range. Cross-correlation methods can be used to reduce the DDS noise. Instruments using those techniques are available that automatically make both time and frequency domain measurements.

9.4. Measurement Problems and Pitfalls

It can be difficult to distinguish between a bad unit under test and a bad measurement. When problems occur in time-domain frequency stability measurements, they usually cause results that are worse than expected. It is nearly impossible for a measurement problem to give better than correct results, and there is considerable justification in saying that the best results are the correct ones. Two possible exceptions to this are (1) misinterpretation of the scale factor, and (2) inadvertent coherency (e.g., injection locking of one source to another due to inadequate isolation. Lack of stationarity (changes in the source itself), while not a measurement problem per se, must also be considered. In general, the more devices available and the more measurements being made, the easier it is to sort things out.

One common problem is hum that contaminates the measurements due to ground loops. Because the measurement interval is usually much longer than the period of the power line frequency, and not necessarily coherent with it, aliased "beats" occur in the frequency record. Inspection of the raw data can show this, and the best cure is often isolation transformers in the signal leads. In fact, this is a wise precaution to take in all cases.

All sorts of other mechanisms (electrical, electromagnetic, magnetic, thermal, barometric, vibrational, acoustic, etc.) exist that can interfere with time domain frequency measurements. Think about all the environmental sensitivities of the unit under test, and control as many of them as possible. Be alert to day-night and weekly cycles that indicate human interference. Stories abound about correlations between elevators moving and cars coming and going ("auto-correlations") that have affected clock measurements. Think about what can have changed in the overall test setup. Slow periodic fluctuations will show up more distinctly in an all tau (rather than an octave tau) stability plot.

In high-precision measurements, where picoseconds matter (e.g. $1 \times 10^{-15} = 1$ ps/1000 seconds), it is important to consider the mechanical rigidity of the test setup (e.g. 1 ps = 0.3 mm). This includes the electrical length (phase stability) of the connecting cables. Teflon dielectric is an especially bad choice for temperature stability, while foamed polyethylene is much better. Even a few degrees of temperature variation will cause the phase of a high-stability source to "breathe" as it passes through 100 ft of coaxial cable.

Phase jumps are a problem that should never be ignored. Examination of the raw phase record is critical because a phase jump (frequency impulse) appears strangely in a frequency record as a white FM noise characteristic [10]. Some large phase jumps are related to the carrier period (e.g., a malfunctioning digital frequency divider).

It is difficult to maintain the integrity of a measuring system over a long period, but, as long as the operating conditions of the unit under test and the reference are undisturbed, gaps in the data record may be acceptable. An uninterruptible power system is indispensable to maintain the continuity of a long run.

9.5. Measuring System Summary

A comparison of the relative advantages and disadvantages of these methods is shown in the following table:

Table 21. Comparison of time and frequency measurement methods.

Method	Advantages	Disadvantages
Divider & time interval counter	Provides phase data	Modest resolution
	Covers wide range of carrier frequencies	
	Easily expandable at low cost	Not suitable for short tau
Mixer and period counter	Resolution enhanced by heterodyne factor	No phase data
	Provides direct frequency data	No frequency sense
	Usable for short tau	Requires offset reference
	Expandable at reasonable cost	Single carrier frequency
Dual mixer time difference	High resolution, low noise	Single carrier frequency
	Provides phase data	
	Offset reference noise & inaccuracy cancels	Relatively complex
	No fixed reference channel	

It is preferable to make continuous zero-dead-time phase measurements at regular intervals, and a system using a dual-mixer time interval measurement is recommended. An automated high-resolution multi-channel clock (phase) measuring system with a high-performance (e.g., hydrogen maser) reference is a major investment, but one that can pay off in better productivity. It is desirable that the measurement control, data storage, and analysis functions be separated to provide robustness and networked access to the data. A low-noise reference not only supports better measurement precision but also allows measurements to be made faster (with less averaging).

9.6. Data Format

A one-column vector is all that is required for a phase or frequency data array. Because the data points are equally spaced, no time tags are necessary. Nevertheless, the use of time tags is recommended (see section 9.4 below), particularly to identify anomalies or to compare several sources. Time tagging is generally required for archival storage of clock measurements, but a single vector of extracted gap-filled data is sufficient for analysis. The recommended unit for phase data is seconds, while frequency data should be in the form of dimensionless fractional frequency. Double-precision exponential ASCII numeric format is recommended for ease of reading into most analysis software, with comma or space-delimited fields and one data point per line. The inclusion of comments and headers can pose problems, but most software will reject lines that start with a "#" or some other non-numeric character.

9.7. Data Quantization

The phase or frequency data must be gathered with sufficient resolution to show the variations of interest, and it must be represented with sufficient precision to convey those variations after removal of fixed offsets (see section 10.1 below). Nevertheless, highly quantized data can still contain useful information, especially after they are combined into longer averaging times. An example of highly quantized frequency data is the random telegraph signal shown below. Although these data have a non-Gaussian amplitude distribution (their histogram consists of two spikes), the random occurrences of the two levels produce a white FM noise characteristic.
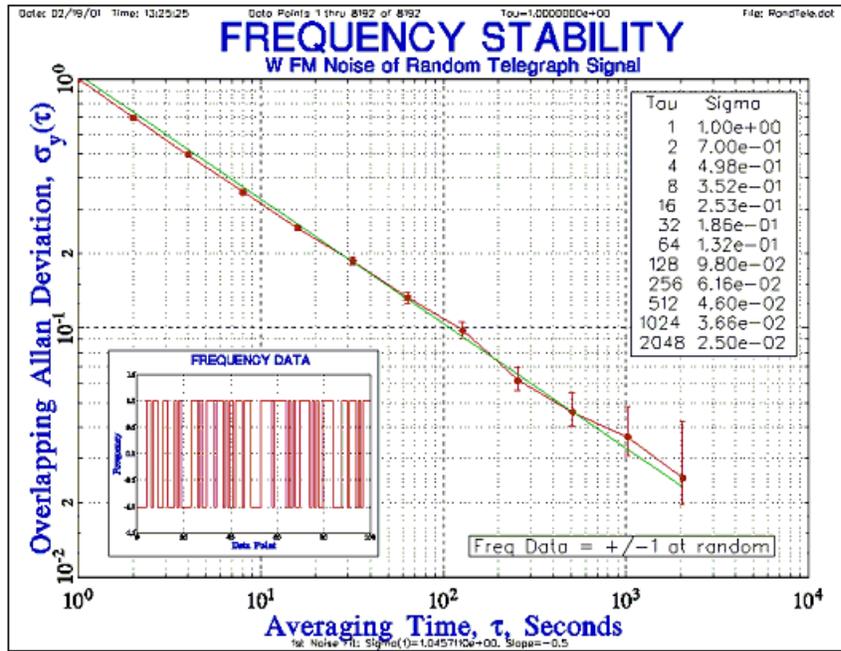
Figure 37. Random telegraph signal as an example of highly quantized frequency data.

9.8. Time tags

Time tags are often associated with phase or frequency data, and can be usefully applied to the analysis of these data.

Time tags are highly desirable for frequency stability measurements, particularly for identifying the exact time of some anomaly. The preferred time tag is the Modified Julian Date (MJD) expressed as a decimal fraction and referenced to UTC. Based on the astronomical Julian Date, the number of days since noon on January 1, 4713 BC, the MJD is the Julian Date – 2 4000 000.5. It is widely used, purely numeric, can have any required resolution, is easily converted to other formats, is non-ambiguous over a two-century (1900 to 2099) range, and is free from seasonal (daylight saving time) discontinuities

Analysis software can easily convert the MJD into other formats such as year, month, day, hour, minute, and second. The MJD (including the fractional day) can be obtained from the C language `time()` function by dividing its return value by 86 400 and adding 40 587.

9.9. Archiving and Access

There is no standard way to archive and access clock data. For some purposes, it is sufficient to simply save the raw phase or frequency data to a file, identifying it only by the file name. At the other extreme, multichannel clock measuring systems may require an elaborate database to store a large collection of data, keep track of the clock identities and transactions, provide security and robust data integrity, and serve the archived data via a network. It may also be necessary to integrate the clock data with other information (e.g., temperature) from a data acquisition system.

References for Measuring Systems

1. D.W. Allan, "Picosecond Time Difference Measurement System," *Proc. 29th Annu. Symp. on Freq. Contrl.*, pp. 404-411, May 1975.
2. D.A. Howe, D.W. Allan and J.A. Barnes, "Properties of Signal Sources and Measurement Methods," *Proc. 35th Annu. Symp. on Freq. Contrl.*, May 1981, pp. 1-47.
3. S.R. Stein, "Frequency and Time - Their Measurement and Characterization," Chapter 12, pp.191-416, Precision Frequency Control, Vol. 2, Edited by E.A. Gerber and A. Ballato, Academic Press, New York, 1985, ISBN 0-12-280602-6.
4. S. Stein, D.Glaze, J. Levine, J. Gray, D. Hilliard, D. Howe and L Erb, "Performance of an Automated High Accuracy Phase Measurement System," *Proc. 36th Annu. Freq. Contrl. Symp*., June 1982, pp. 314-320.
5. S.R. Stein and G.A. Gifford, "Software for Two Automated Time Measurement Systems," *Proc. 38th Annu. Freq. Contrl.* Symp, June 1984, pp. 483-486.
6. Data Sheet, Model 5110A Time Interval Analyzer, Symmetricom, Inc. (formerly Timing Solutions Corp.), Boulder, CO 80301 USA.
7. Data Sheet, A7 Frequency & Phase Comparator Measurment System, Quartzlock (UK) Ltd., Gothic, Plymouth Road, Totnes, Devon, TQ9 5LH England.
8. C.A. Greenhall, "Oscillator-Stability Analyzer Based on a Time-Tag Counter," *NASA Tech Briefs*, NPO-20749, May 2001, p. 48.
9. Gernot M.R. Winkler, "Introduction to Robust Statistics and Data Filtering," Tutorial at 1993 IEEE Freq. Contrl. Symp., Sessions 3D and 4D, June 1, 1993.
10. C.A. Greenhall, "Frequency Stability Review," *Telecommunications and Data Acquisition Progress Report 42-88*, Oct-Dec 1986, Jet Propulsion Laboratory, Pasadena, CA, pp. 200-212, Feb. 1987.
11. Data Sheet, Model 5120A Phase Noise Test Set, Symmetricom, Inc. (formerly Timing Solutions Corp.), Boulder, CO 80301 USA.

10   Analysis Procedure

A frequency stability analysis can proceed along several paths, as the circumstances dictate. Nevertheless, the following example shows a typical analysis flow. Using simulated data for a high-stability rubidium frequency standard, the purpose of the analysis is to characterize the noise in the presence of an outlier, large frequency offset and significant drift.

Table 22. An example of a typical analysis flow.

Step	Description	Plot

1	Open and examine a phase data file. The phase data is just a ramp with slope corresponding to frequency offset.
	
	`0.000000000000000e+00`
	`8.999873025741449e-07`
	`1.799890526185009e-06`
	`2.699869215098003e-06`
	`3.599873851209537e-06`
	`4.499887627663997e-06`
	`5.399836191440859e-06`
	`6.299833612216789e-06`
	`7.199836723454638e-06`
	`8.099785679257264e-06`
	`8.999774896024524e-06`
	`9.899732698242008e-06`
	　　　　`etc.`	
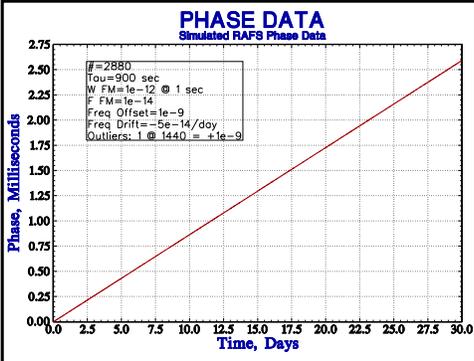
2	Convert the phase data to frequency data and examine it. An obvious outlier exists that must be removed to continue the analysis. **Visual inspection of data is an important preprocessing step!** Analyst judgment may be needed for less obvious outliers.	

3	In an actual analysis, one should try to determine the cause of the outlier. The frequency spike of 1×10^{-9} corresponds to a phase step of 900 ns over a single 900-s measurement interval, nine 10 MHz carrier cycles. Data taken at a higher rate would help to determine whether the anomaly happened instantaneously or over some finite period. Timetags can help to relate the outlier to other external events.	
3	Remove the outlier. The noise and drift are now visible. A line shows a linear fit to the frequency data, which appears to be quite appropriate.	
4	Remove the frequency offset from the phase data. The resulting quadratic shape is due to the frequency drift. One can just begin to see phase fluctuations around the quadratic fit to the phase data.	
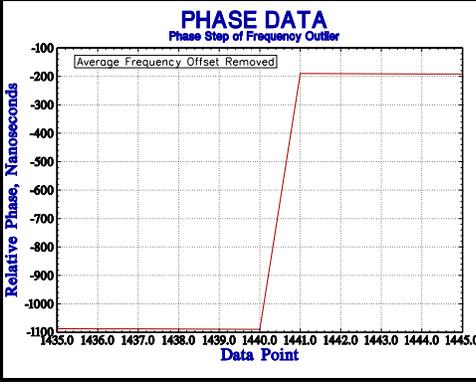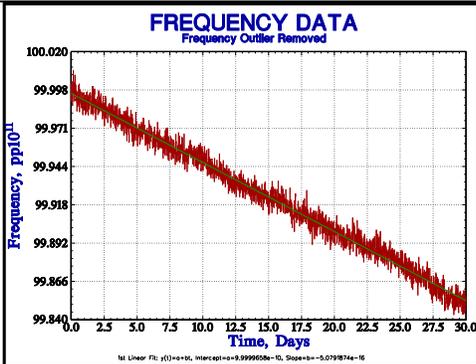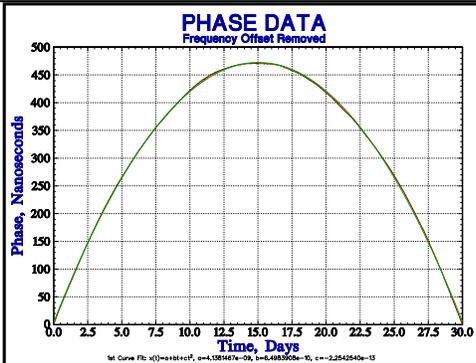
5	Remove the frequency drift, leaving the phase residuals for noise analysis, which is now clearly visible. Some experience is needed to interpret phase data like these. Remember that frequency corresponds to the slope of the phase, so the frequency is lowest near the end of the record, where the phase slope is the most negative.	
6	Convert the phase residuals to frequency residuals. Alternatively, remove the frequency drift from the frequency data of Step #3. There are subtle differences in removing the linear frequency drift as a quadratic fit to the phase data compared with removing it as a linear fit to the frequency data (different noise models apply). Other drift models may be more appropriate. Analyst judgment is needed to make the best choices.	
7	Perform a stability analysis using the overlapping Allan deviation. The results show white FM noise at short averaging times ($\tau^{-0.5}$ slope), and flicker FM noise at longer tau (t^0 slope), both at the simulated levels shown in the annotations of the first plot.	

10.1. Data Precision

There are relatively few numerical precision issues relating to the analysis of frequency stability data. One exception, however, is phase data for a highly stable frequency source having a relatively large frequency offset. The raw phase data will be essentially a straight line (representing the frequency offset), and the instability information is contained in the small deviations from the line. A large number of digits must be used unless the frequency offset is removed by subtracting a linear term from the raw phase data. Similar considerations apply to the quadratic phase term (linear frequency drift). Many frequency stability measures involve averages of first or second differences. Thus, while their numerical precision obviously depends upon the variable digits of the data set, there is little error propagation in forming the summary statistics.

10.2. Preprocessing

Preprocessing of the measurement data is often necessary before the actual analysis is performed, which may require data averaging, or removal of outliers, frequency offset, and drift.

Phase data may be converted to frequency data, and vice versa. Phase and frequency data can be combined for a longer averaging time. Frequency offset may be removed from phase data by subtracting a line determined by the average of the first differences, or by a least squares linear fit. An offset may be removed from frequency data by normalizing it to have an average value of zero. Frequency drift may be removed from phase data by a least squares or three-point quadratic fit, or by subtracting the average of the second differences. Frequency drift may be removed from frequency data by subtracting a least-squares linear fit, by subtracting a line determined by the first differences or by calculating the drift from the difference between the two halves of the data. The latter, called the bisection drift, is equivalent to the three-point fit for phase data. Other more specialized log and diffusion models may also be used. The latter are particularly useful to describe the stabilization of a frequency source. In general, the objective is to remove as much of the deterministic behavior as possible, obtaining random residuals for subsequent noise analysis.

10.3. Gaps, Jumps, and Outliers

It is common to have gaps and outliers in a set of raw frequency stability data. Missing or erroneous data may occur due to power outages, equipment malfunctions, and interference. For long-term tests, it may not be possible or practical to repeat the run, or otherwise avoid such bad data points. Usually the reason for the gap or outlier is known, and it is particularly important to explain all phase discontinuities. Plotting the data will often show the bad points, which may have to be removed before doing an analysis to obtain meaningful results.

> Gaps, jumps and outliers can occur in frequency measurements and they must be handled before performing a stability analysis. Methods are available to fill gaps and to correct for outliers in a consistent manner.

Frequency outliers are found by comparing each data point with the median value of the data set plus or minus some multiple of the absolute median deviation. These median statistics are more robust because they are insensitive to the size of the outliers. Outliers can be replaced by gaps or filled with interpolated values.

Frequency jumps can also be a problem for stability analysis. Their occurrence indicates that the statistics are not stationary, and it may be necessary to divide the data into portions and analyze them separately.

Gaps and outliers can occur in clock data due to problems with the measuring system or the frequency source itself. Like death and taxes, gaps and outliers can be avoided but not eliminated.

10.4. Gap Handling

Gaps should be included to maintain the proper implied time interval between measurements, and a value of zero (0) is often used to denote a gap. For phase data, zero should be treated as valid data if it is the first or last point. For fractional frequency data, valid data having a value of zero can be replaced by some very small value (e.g., 1e−99). Many analysis functions can produce meaningful results for data having gaps by simply skipping those points that involve a gap. For example, in the calculation of the Allan variance for frequency data, if either of the two points involved in the first difference is a gap, that Allan variance pair is skipped in the summation.

Gaps may be filled in phase or frequency data by replacing them with interpolated values, by first removing any leading and trailing gaps, and then using the two values immediately before and after any interior gaps to determine linearly interpolated values within the gap.

A zero value in fractional frequency data can also occur as the result of the conversion of two equal adjacent phase data points (perhaps because of limited measurement resolution), and the value should be adjusted to, say, 1e−99 to distinguish it from a gap.

10.5. Uneven Spacing

Unevenly spaced phase data can be handled if they have associated time tags by using the individual time tag spacing when converting it to frequency data. Then, if the tau differences are reasonably small, the data may be analyzed by using the average time tag spacing as the analysis tau, in effect placing the frequency data on an average uniform grid. While completely random data spacing is not amenable to this process, tau variations of ±10 % will yield reasonable results as long as the exact intervals are used for the phase to frequency conversion.

10.6. Analysis of Data with Gaps

Care must be taken when analyzing the stability of data with missing points and/or gaps. Missing points can be found by examining the time tags associated with the data, and gaps can then be inserted as placeholders to maintain equally spaced data. Similarly, outliers can be replaced with gaps for the same reason. These gaps can span multiple points. Some analysis processes can be performed with data having gaps by skipping over them, perhaps at some speed penalty, but other calculations cannot be. It is therefore often necessary to replace the gaps with interpolated values. Those points are not real data, however, and, if there are many of them, the results will be suspect. In these cases, judgment is needed to assure a credible result. It may be more prudent to simply analyze a gap-free portion of the data.

10.7. Phase-Frequency Conversions

Phase to frequency conversion is straightforward for data having gaps. Because two phase points are needed to determine each frequency point (as the difference between the phase values divided by their tau), a single phase gap will cause two frequency gaps, and a gap of N phase points causes N + 1 frequency gaps.

Conversion from frequency to phase is more problematic because of the need to integrate the frequency data. The average frequency value is used to calculate the phase during the gap, which can cause a discontinuity in the phase record. Analysis of phase data resulting from the conversion of frequency data having a large gap is not recommended.

10.8. Drift Analysis

Drift analysis functions generally perform well for data having gaps, provided that missing data are represented by gaps to maintain a regular time sequence.

10.9. Variance Analysis

Variance analysis functions can include provisions for handling gaps. Some of these functions yield satisfactory results in all cases, while others have speed limitations, or provide unsatisfactory results for data having large gaps. The latter is most apparent at longer averaging times where the averaging factor is comparable to the size of the gap. The speed limitations are caused by more complex gap checking and frequency averaging algorithms, while the poor results are associated with the total variances for which conversion to phase data is required. In all cases, the results will depend on coding details included in addition to the basic variance algorithm. Filling gaps can often help for the total variances. Two general rules apply for the variance analysis of data having large gaps: (1) use unconverted phase data, and (2) check the results against the normal Allan deviation (which has the simplest, fastest gap handling ability).

10.10. Spectral Analysis

Gap filling in spectral analysis functions can affect the low frequency portion of the spectrum.

10.11. Outlier Recognition

The median absolute deviation (MAD) is recommended as its means of outlier recognition. The MAD is a robust statistic based on the median of the data. It is the median of the scaled absolute deviations of the data points from their median value, defined as MAD = Median { | $y(i)$ − m | / 0.6745 }, where m = Median { $y(i)$ }, and the factor 0.6745 makes the MAD equal to the standard deviation for normally distributed data. Each frequency data point, $y(i)$, is compared with the median value of the data set, m, plus or minus the desired multiple of the MAD.

While the definition of an outlier is somewhat a matter of judgment, it is important to find and remove such points in order to use the rest of the data, based on their deviation from the median of the data, using a deviation limit in terms of the median absolute deviation (a 5-sigma limit is common). This is a robust way to determine an outlier, which is then replaced by a gap. An automatic outlier removal algorithm can iteratively apply this method to remove all outliers, which should be an adjunct to, and not a substitute for, visual inspection of the data.

It is important to explain all outliers, thereby determining whether they are due to the measurement process or the device under test. An important first step is to correlate the bad point with any external events (e.g., power outages, equipment failures, etc.) that could account for the problem. Failures of the measurement system, frequency reference, or environmental control are often easier to identify if multiple devices are under test. Obviously, a common gap in all measurement channels points to a failure of the measurement system, while a common change in all measurement readings points to a reference problem. Auxiliary information such as monitor data can be a big help in determining the cause of outliers. A log of all measurement system events should be kept to facilitate outlier identification. Discrete phase jumps are a particular concern, and, if they are related to the RF carrier frequency, may indicate a missing cycle or a problem with a digital divider. A phase jump will correspond to a frequency spike with a magnitude equal to the phase change divided by the measurement interval. Such a frequency spike will produce a stability record that appears to have a (large magnitude) white FM noise characteristic, which can be a source of confusion.

References for Gaps, Jumps, and Outliers

1. D.B. Percival, "Use of Robust Statistical Techniques in Time Scale Formation," Preliminary Report, U.S. Naval Observatory Contract No. N70092-82-M-0579, 1982.
2. Gernot M.R. Winkler, "Introduction to Robust Statistics and Data Filtering," Tutorial at 1993 IEEE Freq. Contrl. Symp., Sessions 3D and 4D, June 1, 1993.
3. V. Barnett and T. Lewis, <u>Outliers in Statistical Data</u>, 3rd Edition, John Wiley & Sons, Chichester, 1994, ISBN 0-471-93094-6.

10.12.   Data Plotting

Data plotting is often the most important step in the analysis of frequency stability. Visual inspection can provide vital insight into the results, and is an important preprocessor before numerical analysis. A plot also shows much about the validity of a curve fit.

Phase data are plotted as line segments connecting the data points. This presentation properly conveys the integral nature of the phase data. Frequency data are plotted as horizontal lines between the frequency data points. This shows the averaging time associated with the frequency measurement, and mimics the analog record from a frequency counter. As the density of the data points increases, there is essentially no difference between the two plotting methods. Missing data points are shown as gaps without lines connecting the adjacent points.

10.13.   Variance Selection

It is the user's responsibility to select an appropriate variance for the stability analysis. The overlapping Allan variance is recommended in most cases, especially where the frequency drift is small or has been removed. The Allan variance is the most widely used time-domain stability measure, and the overlapping form provides better confidence than the original "normal" version. The total and ThêoH variance can be used for even better confidence at large averaging factors (but at the expense of longer computation time). The modified Allan variance is recommended to distinguish between white and flicker PM noise, and, again, a total form of it is available for better confidence at long tau. The time variance provides a good measure of the time dispersion of a clock due to noise, while MTIE measures the peak time deviations. TIE rms can also be used to assess clock performance, but TVAR is generally preferred. Finally, the overlapping Hadamard variance is recommended over its normal form for analyzing stability in the presence of divergent noise or frequency drift. In all cases, the results are reported in terms of the deviations.

The choice of tau interval depends mainly on whether interference mechanisms are suspected that cause the stability to vary periodically. Normally, octave or decade spacing is used (the former has even spacing on a log-log plot, while the latter provides tau multiples of ten). The all tau option can be useful as a form of spectral analysis to detect cyclic disturbances (such as temperature cycling).

10.14.   Three-Cornered Hat

Any frequency stability measurement includes noise contributions from both the device under test and the reference. Ideally, the reference noise would be low enough that its contribution to the measurement is negligible. Or, if the noise of the reference is known, it can be removed by subtracting its variance. A special case is that of two identical units where half of the measured variance comes from each, and the measured deviation can be corrected for one unit by dividing it by √2. Otherwise, it may be useful to employ the so-called "three-cornered hat" method for determining the variance of an individual source. Given a set of three pairs of measurements for three independent frequency sources a, b, and c whose variances

> It is sometimes necessary to determine the individual noise contributions of the two sources that contribute to a variance measurement. Methods exist for doing so by using the results of multiple measurements to estimate the variance of each source.

add

$$\sigma^2_{ab} = \sigma^2_a + \sigma^2_b$$
$$\sigma^2_{ac} = \sigma^2_a + \sigma^2_c ,$$
$$\sigma^2_{bc} = \sigma^2_b + \sigma^2_c$$

(68)

The individual variances may be determined by the expressions

$$\sigma^2_a = \frac{1}{2}\left[\sigma^2_{ab} + \sigma^2_{ac} - \sigma^2_{bc}\right]$$

$$\sigma^2_b = \frac{1}{2}\left[\sigma^2_{ab} + \sigma^2_{bc} - \sigma^2_{ac}\right],$$

(69)

$$\sigma^2_c = \frac{1}{2}\left[\sigma^2_{ac} + \sigma^2_{bc} - \sigma^2_{ab}\right]$$

Although useful for determining the individual stabilities of units having similar performance, the method may fail by producing negative variances for units that have widely differing stabilities, if the units are correlated, or for which there are insufficient data. The three sets of stability data should be measured simultaneously. The three-cornered hat method should be used with discretion, and it is not a substitute for a low noise reference. It is best used for units having similar stability (e.g., to determine which unit is best). Negative variances are a sign that the method is failing (because it was based on insufficient measurement data, or because the units under test have disparate or correlated stability). This problem is most likely to arise at long tau.

The three-cornered hat function may be used to correct a stability measurement for the noise contribution of the reference, as shown in Figure 38.
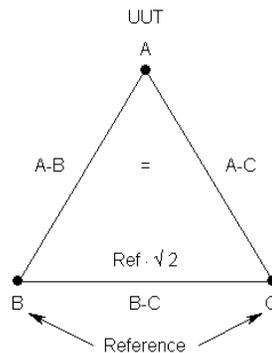
Figure 38. Three-cornered hat function.

The Unit Under Test (UUT), denoted as source A, is measured against the reference, denoted by B and C, by identical stability data files A-B and A-C. The reference is measured against itself by stability data file B-C, which contains the a priori reference stability values multiplied by $\sqrt{2}$.

An example of the use of the three-cornered hat function to correct stability data for reference noise is shown below. Simulated overlapping Allan deviation stability data for the unit under test versus the reference were created by generating and analyzing 512 points of frequency data with tau = 1 s and $\sigma_y(1\ s)$ = 1e−11. The resulting stability data are shown in the following table.

Table 23. Stability data for unit under test versus reference.

Tau	#	Sigma	Min Sigma	Max Sigma
1.000e+00	511	9.448e-12	9.108e-12	9.830e-12
2.000e+00	509	7.203e-12	6.923e-12	7.520e-12
4.000e+00	505	5.075e-12	4.826e-12	5.367e-12
8.000e+00	497	3.275e-12	3.058e-12	3.546e-12
1.600e+01	481	2.370e-12	2.157e-12	2.663e-12
3.200e+01	449	1.854e-12	1.720e-12	2.025e-12
6.400e+01	385	1.269e-12	1.147e-12	1.441e-12
1.280e+02	257	5.625e-13	4.820e-13	7.039e-13

A similar stability file is used for the reference. Since it represents a measurement of the reference against itself, the Allan deviations of the reference source are multiplied by $\sqrt{2}$. Simulated overlapping Allan deviation stability data for the reference versus the reference was created by generating 512 points of frequency data with tau = 1 second and $\sigma_y(1) = 1.414e{-}12$.

Table 24. Stability data for unit under test versus reference.

Tau	#	Sigma	Min Sigma	Max Sigma
1.000e+00	511	1.490e-12	1.436e-12	1.550e-12
2.000e+00	509	1.025e-12	9.854e-13	1.070e-12
4.000e+00	505	7.631e-13	7.257e-13	8.070e-13
8.000e+00	497	5.846e-13	5.458e-13	6.329e-13
1.600e+01	481	3.681e-13	3.349e-13	4.135e-13
3.200e+01	449	2.451e-13	2.152e-13	2.924e-13
6.400e+01	385	1.637e-13	1.368e-13	2.173e-13
1.280e+02	257	1.360e-13	1.058e-13	2.285e-13
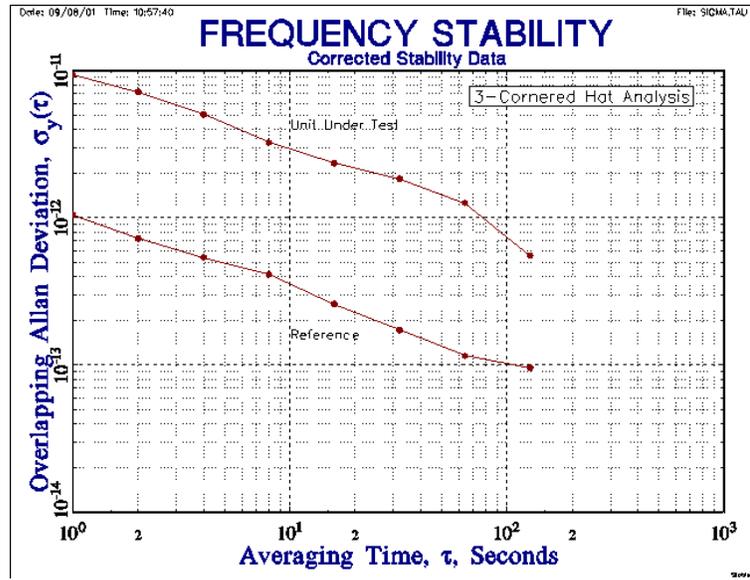
Figure 39. Corrected UUT and reference stabilities.

Here the reference stability is about $1 \times 10^{-12}\tau^{-1/2}$, and the corrected UUT instability is slightly less than the uncorrected values. Note that the B and C columns of corrected stability values both represent the reference source. Appropriate use of the three-cornered hat method to correct stability measurements for reference noise applies where the reference stability is between three to 10 times better than that of the unit under test. The correction is negligible for more the latter (see above), and has questionable confidence for less than the former (and a better reference should be used).

The error bars of the individual variances may be set using χ^2 statistics by first determining the reduced number of degrees of freedom associated with the three-cornered hat process [11, 12]. The fraction of remaining degrees of freedom for unit i as a result of performing a three-cornered hat instead of measuring against a perfect reference is given by:

$$\Gamma = \frac{2 \cdot \sigma_i^4}{2 \cdot \sigma_i^4 + \sigma_a^2 \cdot \sigma_b^2 + \sigma_a^2 \cdot \sigma_c^2 + \sigma_b^2 \cdot \sigma_c^2} \ . \tag{70}$$

The ratio of the number of degrees of freedom is 0.4 for three units having the same stability, independent of the averaging time and noise type.

The three-cornered hat technique can be extended to M clocks (subject to the same restriction against negative variances) by using the expression

$$\sigma_i^2 = \frac{1}{M-2}\left[\sum_{j=1}^{M} \sigma_{ij}^2 - B \right], \tag{71}$$

where

95

$$B = \frac{1}{2(M-1)} \sum_{k=1}^{M} \left[\sum_{j=1}^{M} \sigma_{kj}^2 \right], \tag{72}$$

and the σ^2_{ij} are the measured Allan variances for clock i versus j at averaging time τ. Using $\sigma^2_{ii} = 0$ and $\sigma^2_{ij} = \sigma^2_{ji}$, we can easily write closed-form expressions for the separated variances from measurements of M clocks.

References for Three-Cornered Hat

1. The term "Three-cornered hat" was coined by J.E. Gray of NIST.
2. J.E. Gray and D.W. Allan, "A Method for Estimating the Frequency Stability of an Individual Oscillator," *Proceedings of the 28th Annual Symposium on Frequency Control*, May 1974, pp. 243-246.
3. S.R. Stein, "Frequency and Time - Their Measurement and Characterization," Chapter 12, Section 12.1.9, Separating the Variances of the Oscillator and the Reference, pp. 216-217, *Precision Frequency Control* , Vol. 2, Edited by E.A. Gerber and A. Ballato, Academic Press, New York, 1985, ISBN 0-12-280602-6.
4. J. Groslambert, D. Fest, M. Oliver and J.J. Gagnepain, "Characterization of Frequency Fluctuations by Crosscorrelations and by Using Three or More Oscillators," *Proceedings of the 35th Annual Frequency Control Symposium*, May 1981, pp. 458-462.
5. P. Tavella and A. Premoli, "A Revisited Tree-Cornered Hat Method for Estimating Frequency Standard Instability," *IEEE Transactions on Instrumentation and Measurement*, IM-42, February 1993, pp. 7-13.
6. P. Tavella and A. Premoli, "Characterization of Frequency Standard Instability by Estimation of their Covariance Matrix," *Proceedings of the 23rd Annual Precise Time and Time Interval (PTTI) Applications and Planning Meeting*, December 1991, pp. 265-276.
7. P. Tavella and A. Premoli, "Estimation of Instabilities of N Clocks by Measuring Differences of their Readings," *Metrologia*, Vol**.** 30, No. 5, 1993, pp. 479-486.
8. F. Torcaso, C.R. Ekstrom, E.A. Burt and D.N. Matsakis, "Estimating Frequency Stability and Cross-Correlations," *Proceedings of the 30th Annual Precise Time and Time Interval (PTTI) Systems and Applications Meeting*, December 1998, pp. 69-82.
9. F. Torcaso, C. Ekstrom, E. Burt and D. Matsakis, "Estimating the stability of N clocks with correlations," *IEEE Trans. Instrum. Meas. 47(5)*, 1183, 2000.
10. C. Audoin and B. Guinot, *The Measurement of Time,* Cambridge University Press, ISBN 0-521-00397-0, 2001, Section 5.2.8.
11. C. Ekstrom and P. Koppang, "Three-Cornered Hats and Degrees of Freedom," *Proceedings of the 33rd Annual Precise Time and Time Interval (PTTI) Systems and Applications Meeting*, November 2001, pp. 425-430.
12. C.R. Ekstrom and P.A. Koppang, "Error Bars for Three-Cornered Hats," *IEEE Trans. UFFC*, Vol. 53, No. 5, pp. 876-879, May 2006.
13. J.A. Barnes, *The Analysis of Frequency and Time Data*, Austron, Inc. May 1992.

10.15.   Reporting

The results of a stability analysis are usually presented as a combination of textual, tabular, and graphic forms. The text describes the device under test, the test setup, and the methodology of the data preprocessing and analysis, and summarizes the results. The assumptions and analysis choices should be documented so that the results could be reproduced. The report often includes a table of the stability statistics. Graphical presentation of the data at each stage of the analysis is generally the most important aspect of presenting the results. For example, these are often a series of plots showing the phase and frequency data with an aging fit, phase and frequency residuals with the aging removed, and stability plots with noise fits and error bars. Plot titles, subtitles, annotations and inserts can be used to clarify and emphasize the data presentation. The results of several stability runs can be combined, possibly along with specification limits, into a single composite plot. The various elements can be combined into a single electronic document for easy printing and transmittal.

11 Case Studies

This section contains several case studies to further illustrate methodologies of frequency stability analysis.

11.1. Flicker Floor of a Cesium Frequency Standard

The purpose of this analysis is to examine the flicker floor of a commercial cesium beam tube frequency standard. An instrument of this kind can be expected to display a white FM noise characteristic out to long averaging times. At some point, however, the unit will typically "flicker out", with its stability plot flattening out to a flicker FM noise characteristic, beyond which its stability will no longer improve. Determining this point requires a lengthy and expensive run, and it is worthwhile to use an analytical method that provides the best information at long averaging factors. The effort of a more elaborate analysis is far easier than extending the measurement time by weeks or months.

Table 25. Flicker floor of a cesium frequency standard.

This is the five-month frequency record for the unit under test. It has already been "cleaned up" for any missing points, putting in gaps as required to provide a time-continuous data set. The data look very "white" (see Section 6.1), the frequency offset is very small ($+7.2 \times 10^{-14}$), and there is no apparent drift ($+1.4 \times 10^{-16}$/day). Overall, this appears to be a very good record.	
An overlapping Allan deviation plot shows a white FM noise level of about $8.2 \times 10^{-12}\tau^{-1/2}$ out to about 5 days, where the stability levels off at about 1.2×10^{-14}. While this is very respectable behavior, one wonders what the stability actually is at the longer averaging times. But to gain meaningful confidence using ADEV there, the run would have to be extended by several months.	

97

The total deviation can provide better confidence at the longer averaging times. It seems to indicate that the stability continues to improve past 10 days, where is drops below 1×10^{-14}, but the results are not conclusive.

FREQUENCY STABILITY
Cesium Beam Frequency Standard

Tau	Sigma
1.00e+02	8.12e−13
2.00e+02	5.79e−13
4.00e+02	4.11e−13
8.00e+02	2.91e−13
1.60e+03	2.06e−13
3.20e+03	1.45e−13
6.40e+03	1.02e−13
1.28e+04	7.29e−14
2.56e+04	5.28e−14
5.12e+04	3.72e−14
1.02e+05	2.53e−14
2.05e+05	1.69e−14
4.10e+05	1.24e−14
8.19e+05	1.19e−14
1.64e+06	1.18e−14
3.28e+06	6.95e−15
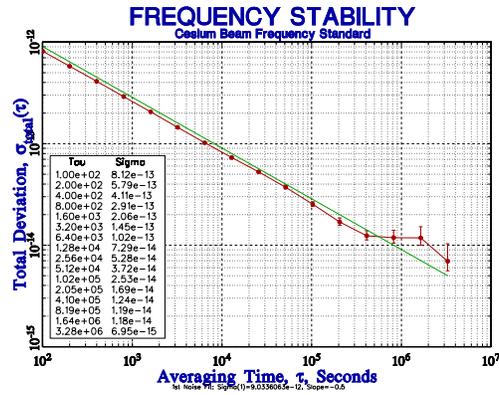
Thêo1 can provide even better long-term information, at the expense of a longer calculation time. This Thêo1 plot seems to show even more clearly that the stability continues to improve at longer averaging times, well into the pp10^{15} range. It assumes white FM noise (no Thêo1 bias removal).

FREQUENCY STABILITY
Cesium Beam Frequency Standard

Tau	Sigma
7.50e+02	3.00e−13
1.50e+03	2.12e−13
3.00e+03	1.50e−13
6.00e+03	1.05e−13
1.20e+04	7.52e−14
2.40e+04	5.37e−14
4.80e+04	3.80e−14
9.60e+04	2.70e−14
1.92e+05	1.83e−14
3.84e+05	1.27e−14
7.68e+05	9.95e−15
1.54e+06	9.48e−15
3.07e+06	7.54e−15
6.14e+06	6.19e−15
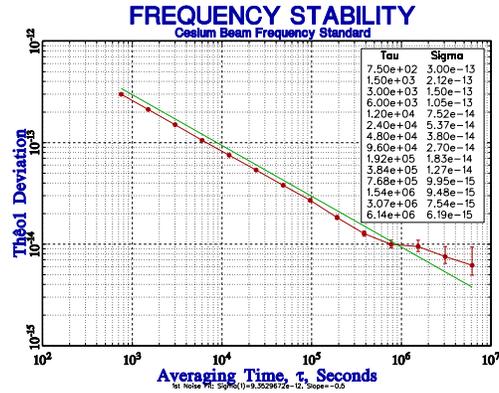
A ThêoH analysis with automatic bias removal combines AVAR in short and mid-term averaging times and ThêoBR in long term. ThêoH does not require any explicit knowledge of the type of noise. It appears that one noise prevails in this data run, in which case Thêo1 (shown above) bias is detected as 1.676, intermediate between white and flicker FM noises for medium to large averaging times. Thus, Thêo1 results are essentially the same as ThêoH.

FREQUENCY STABILITY
Cesium Beam Frequency Standard

Theo1 Bias = 1.676

AF	Sigma	Thêo1
1	3.66e−13	
2	2.61e−13	
4	1.84e−13	
8	1.29e−13	
16	9.22e−14	
32	6.56e−14	
64	4.69e−14	
128	3.35e−14	
256	2.19e−14	
512	1.53e−14	
1024	1.22e−14	
2048	1.23e−14	
4096	1.14e−14	
8192		9.76e−15
16384		8.01e−15
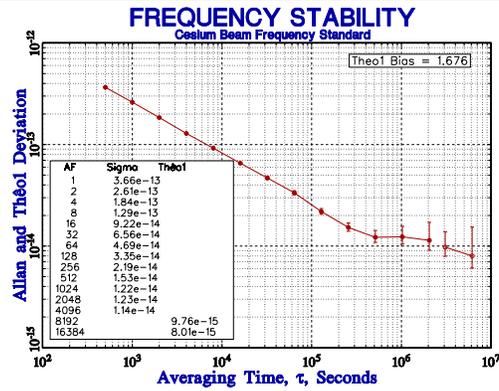
We can conclude that this cesium beam frequency standard reaches a stability slightly better than 1×10^{-14} at an averaging time on the order of 1 month.

11.2. Hadamard Variance of a Source with Drift

Table 26. Hadamard variance of a source with drift.

These frequency data simulate a typical rubidium frequency standard (RFS) with a combination of white and flicker FM noise, plus significant frequency drift.	
If an Allan variance analysis is performed directly on these data without drift removal, the stability at the longer averaging times is degraded. In particular, the stability plot has a τ^{+1} slope beyond 10^5 seconds that corresponds to the drift (i.e., about 1×10^{-12} at 5 days).	
If the linear frequency drift is removed before performing the AVAR stability analysis, the stability plot shows a white FM noise $(\tau^{-1/2})$ characteristic changing to flicker FM noise (t^0) at longer averaging times. It is usually best to use a stability plot only to show the noise, and analyzing and removing the drift separately.	

The Hadamard variance is insensitive to linear frequency drift. It can therefore be used to perform a stability analysis on the original data without first having to remove the drift. The HDEV results are essentially identical to those of the drift-removed ADEV. This can be more convenient, especially when analyzing a combination of sources having differing amounts of drift (e.g., cesium and rubidium units).
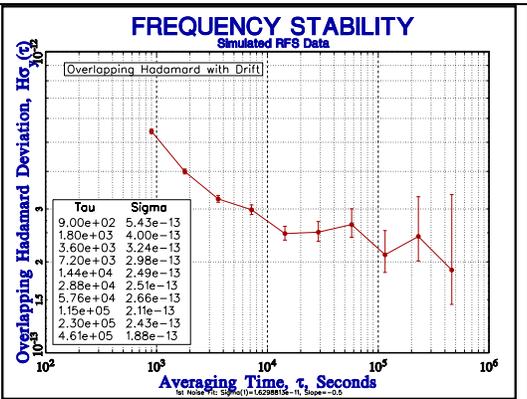
11.3. Phase Noise Identification

Consider the problem of identifying the dominant type of phase noise in a frequency distribution amplifier. Assume that time domain stability measurements have been made comparing the amplifier's input and output signals. How should these data be analyzed to best determine the noise type? Simulated white and flicker PM noise at equal levels are analyzed below in several ways to demonstrate their ability to perform this noise identification.

Table 27. Phase noise identification

White PM	Flicker PM
Examination of the phase data is a good first step. An experienced analyst will immediately notice the difference between the white and flicker noise plots, but will find it harder to quantify them.	
In contrast, the frequency data show little noise type discrimination, because the differencing process whitens both. Examination of the frequency data would be appropriate to distinguish between white and flicker FM noise, however.	

The Allan deviation is not able to distinguish between white and flicker PM noise. Both have slopes of about 1, as shown by the superimposed noise fit lines.
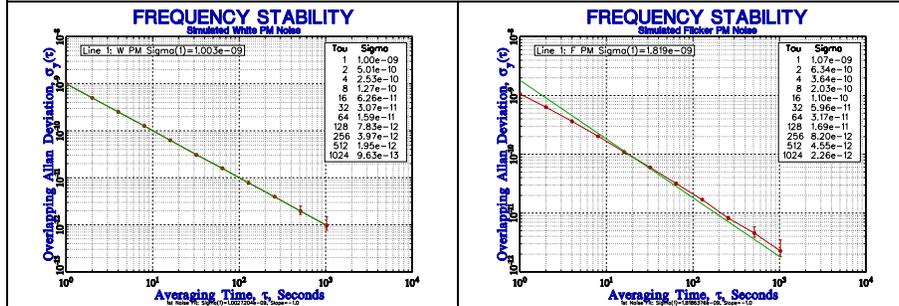
The modified Allan deviation, by virtue of its additional phase averaging, is able to distinguish between white and flicker PM noise, for which the slopes are –1.5 and –1.0, respectively.
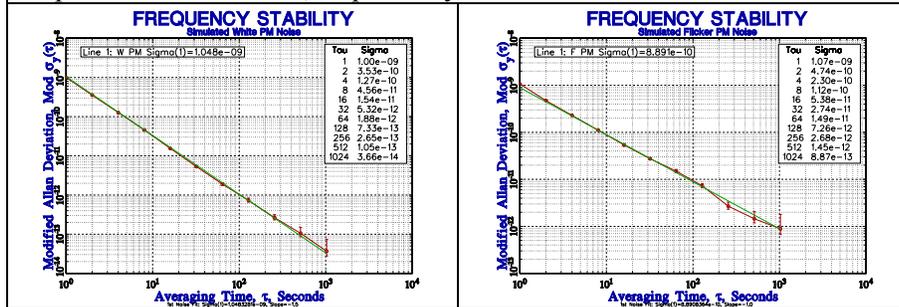
Even better discrimination is possible with the autocorrelation function. The lag 1 ACF is 0.006 and 0.780 for these white and flicker PM noise data, and is able to quantitatively estimate the power law noise exponents as +1.99 and +0.93, respectively. That ID method is quite effective even for mixed noise types.
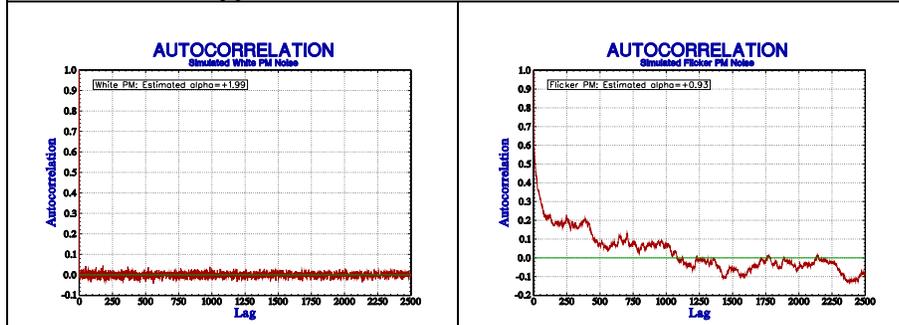
11.4. Detection of Periodic Components

Frequency stability data can contain periodic fluctuations due to external interference, environmental effects, or the source itself, and it can therefore be necessary to detect discrete spectral components when performing a stability analysis. This example uses a set of 50 000 points of $\tau = 1$ second simulated white FM noise at a level of $1 \times 10^{-11}\tau^{-1/2}$ that contains sinusoidal interference with a period of 500 s at a peak level of 1×10^{-12}. The latter simulates interference that might occur as the result of air conditioner cycling. Several analytical methods are used to detect this periodic component.

Table 28. Detection of periodic components.

The interference level is too low to be visible in the full frequency data plot.	
By zooming in, one can see just a hint of the interference (10 cycles over 5000 data points).	
The interference is quite visible in an "all tau" stability plot as a null that occurs first at an averaging time of 500 s (the period of the interference). Here the stability is equal to the underlying white FM noise level.	

The interference is clearly visible in the power spectral density (PSD), which has a bright component at 2 mHz, correspond-ing to the 500 s period of the interference.	**POWER SPECTRUM** Simulated W FM Noise with Periodic Interference Discrete Component with 2mHz Fourier Frequency Multitaper: J=3.0, #=5 Log PSD, $S_y(f)$, Hz^{-1} vs Fourier Frequency, f, Hz
The interference is less visible in an autocorrelation plot, where cyclic behavior is barely noticeable with a 500-lag period. The PSD has equivalent information, but its log scale makes low-level components more apparent.	**AUTOCORRELATION** Simulated W FM Noise with Periodic Interference White FM: Estimated alpha=+0.02 Autocorrelation vs Lag

It is a good analysis policy to examine the power spectral density when periodic fluctuations are visible on a stability plot and periodic interference is suspected.

11.5. Stability Under Vibrational Modulation

This plot shows the stability of an oscillator with a combination of white PM noise and a sinusoidal com-ponent simulating vibrational modula-tion. Nulls in the Allan deviation occur at averaging times equal to the multiples of the 20 Hz sinusoidal mod-ulation period, where the stability is determined by the white PM noise. Peaks in the Allan deviation occur at the modulation half cycles, and have a τ^{-1} envelope set by the vibrational phase modulation.	**FREQUENCY STABILITY** Vibration-Induced Allan Deviation Degradation 1g, 20 Hz vib with 1x10^{-9}/g xtal & 1x10^{-12} W PM noise Overlapping Allan Deviation, $\sigma_y(\tau)$ vs Averaging Time, τ, Seconds

11.6. White FM Noise of a Frequency Spike

The Allan deviation of a frequency record having a large spike (a phase step) has a $\tau^{-1/2}$ characteristic [1]. Thus adding a single large (say 10^6) central outlier to the 1000-point test suite of section 12.4 will give a data set with $\sigma_y(\tau) = [(10^6) \cdot 2/(1000-1)]^{1/2} = 3.16386e4$, as shown in this stability plot.
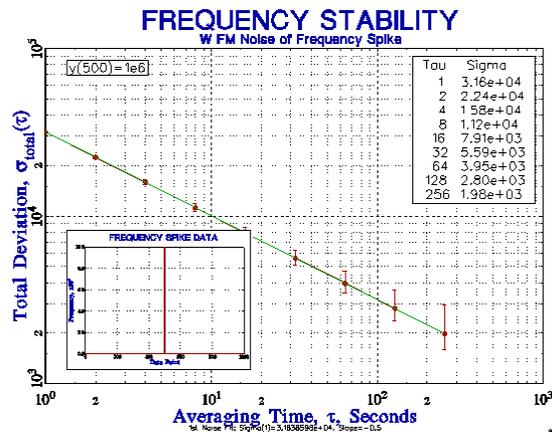
Reference: C.A. Greenhall, "Frequency Stability Review," *Telecommunications and Data Acquisition Progress Report 42-88*, Jet Propulsion Laboratory, Pasadena, CA, Feb. 1987.

11.7. Composite Aging Plots

A composite plot showing the aging of a population of frequency sources can be an effective way to present this information visually, providing a quick comparison of their behaviors. The following figure shows the stabilization period of a production lot of rubidium clocks.

Plots for 40 units are shown with their serial numbers, all plotted with the same scales. In particular, these plots all have a full x-scale time range of nine weeks (1 week/div) and a full y-scale frequency range of 1.5×10^{-11} (1×10^{-12}/div). A diagonal slope downward across the plot corresponds to an aging of about -2.4×10^{-13}/day. All the data have $\tau = 900$ seconds (15 min). A figure like this immediately shows that (a) all the units have negative frequency aging of about the same magnitude that stabilizes in about the same way, (b) there are occasional gaps in some of the records, (c) all of the units have about the same short-term noise, but some of the records are quieter than others in the longer term, (d) some of the units take longer than others to reach a certain aging slope.

The plots, although small, still contain enough detail to allow subtle comparisons very quickly, far better than a set of numbers or bar graphs would do. The eye can easily see the similarities and differences, and can immediately select units based on some criterion, which would be harder to do using a set of larger plots on separate pages. Closer inspection of even these small plots can reveal a lot of quantitative information if one knows the scale factors. Color coding, although not used here, could be used to provide additional information.

These plots are inspired by Edward Tufte's book The Visual Display of Quantitative Information, ISBN 978-0961392109, Graphic Press, 1983.
.
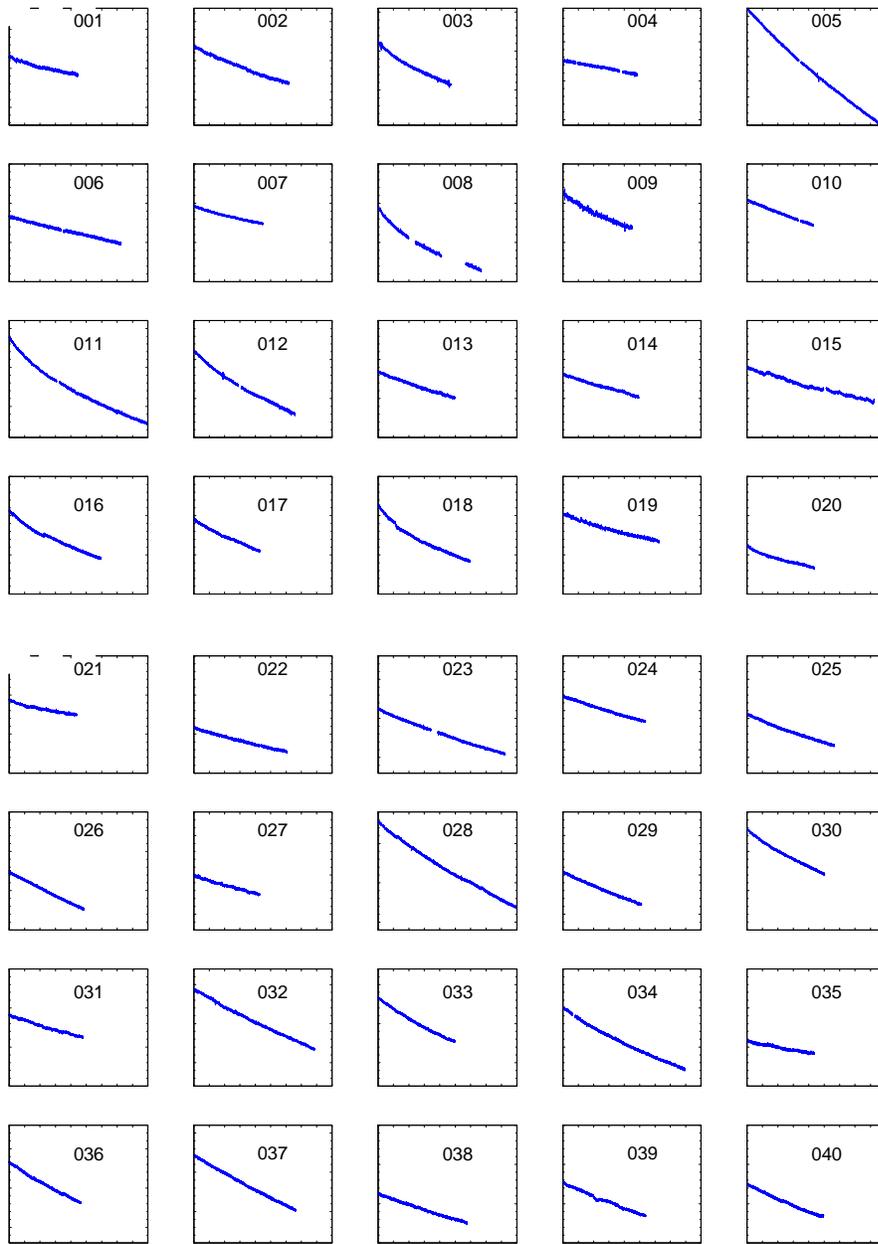
Figure 40. Composite aging plots showing the stabilization period of a production lot of rubidium clocks.

12   Software

Software is necessary to perform a frequency stability analysis, because those calculations generally involve complex specialized algorithms and large data sets needed to interactively perform and document a complete stability analysis. It is convenient to use an integrated software package that combines all of the required

> Specialized software is needed to perform a frequency stability analysis.

analytical tools, operates on current computer hardware and operating systems, includes the latest analytical techniques, and has been validated to produce the correct results.

12.1. Software Validation

Considerable effort is needed to ensure that the results obtained from frequency stability analysis software are correct. Several suggested validation methods are shown. Mature commercially available software should be used whenever possible instead of developing custom software. User feedback and peer review are important. There is a continuing need to validate the custom software used to analyze time domain frequency stability, and the methods listed below can help ensure that correct results are obtained.

Several methods are available to validate frequency stability analysis software:

1. Manual Analysis: The results obtained by manual analysis of small data sets (such as in *NBS Monograph 140* Annex 8.E) can be compared with the new program output. This is always good to do to get a "feel" for the process.
2. Published Results: The results of a published analysis can be compared with the new program output. One important validation method is comparison of the program results against a test suite such as the one in References [1] and [2]. Copies of those test data are available on-line [3].
3. Other Programs: The results obtained from other specialized stability analysis programs (such as that from a previous generation computer or operating system) can be compared with the new program output.
4. General Purpose Programs: The results obtained from industry standard, general purpose mathematical and spreadsheet programs such as MathCAD, Matlab, and Excel can be compared with the new program output.
5. Consistency Checks: The new program should be verified for internal consistency, such as producing the same stability results from phase and frequency data. The standard and normal Allan variances should be approximately equal for white FM noise. The normal and modified Allan variances should be identical for an averaging factor of 1. For other averaging factors, the modified Allan variance should be approximately one-half the normal Allan variance for white FM noise, and the normal and overlapping Allan variances should be approximately equal. The overlapping method provides better confidence of the stability estimates. The various methods of drift removal should yield similar results.
6. Simulated Data: Simulated clock data can also serve as a useful cross check. Known values of frequency offset and drift can be inserted, analyzed, and removed. Known values of power-law noise can be generated, analyzed, plotted, and modeled.

12.2. Test Suites

Tables 29 and 30 summarize the values for several common frequency stability measures for both the classic NBS data set and a 1000-point portable test suite.

12.3. NBS Data Set

A "classic" suite of frequency stability test data is the set of nine 3-digit numbers from Annex 8.E of *NBS Monograph 140* shown in Table 29. Those numbers were used as an early example of an Allan variance calculation. These frequency data is also normalized to zero mean by subtracting the average value, and then integrated to obtain phase values. A listing of the properties of this data set is shown in Table II. While nine data points are not sufficient to calculate large frequency averages, they are, nevertheless, a very useful starting point to verify frequency stability

calculations, since a small data set can easily be entered and analyzed manually. A small data set is also an advantage for finding "off-by-one" errors where an array index or some other integer variable is slightly wrong and hard to detect in a larger data set.

Table 29. NBS Monograph 140, Annex 8.E Test Data

#	Frequency	Normalized frequency	Phase (τ=1)
1	892	103.1111	0.00000
2	809	20.11111	103.11111
3	823	34.11111	123.22222
4	798	9.11111	157.33333
5	671	−117.88889	166.44444
6	644	−144.88889	48.55555
7	883	94.11111	−96.33333
8	903	114.11111	−2.22222
9	677	−111.88889	111.88889
10			0.0000

Table 30. NBS Monograph 140, Annex 8.E Test Data Statistics

Averaging Factor	1	2
# Data Points	9	4
Maximum	903	893.0
Minimum	644	657.5
Average	788.8889	802.875
Median	809	830.5
Linear Slope	−10.20000	−2.55
Intercept	839.8889	809.25
Standard Deviation[a]	100.9770	102.6039
Normal Allan Deviation	91.22945	115.8082
Overlapping Allan Dev	91.22945	85.95287
Modified Allan Dev	91.22945	74.78849
Time Deviation	52.67135	86.35831
Hadamard Deviation	70.80607	116.7980
Overlap Hadamard Dev	70.80607	85.61487
Hadamard Total Dev	70.80607	91.16396
Total Deviation	91.22945	93.90379
Modified Total Dev	75.50203	75.83606
Time Total Deviation	43.59112	87.56794
Note: [a] Sample (not population) standard deviation.		

12.4. 1000-Point Test Suite

The larger frequency data test suite used here consists of 1000 pseudo-random frequency data points. It is produced by the following prime modulus linear congruential random number generator:

$$n_{i+1} = 16807 n_i \text{ Mod } 2147483647 .$$

(73)

This expression produces a series of pseudo-random integers ranging in value from 0 to 2 147 483 646 (the prime modulus, $2^{31}-1$, avoids a collapse to zero). When started with the seed $n_0 = 1\ 234\ 567\ 890$, it produces the sequence $n_1 = 395\ 529\ 916$, $n_2 = 1\ 209\ 410\ 747$, $n_3 = 633\ 705\ 974$, etc. These numbers may be divided by 2 147 483 647 to obtain a set of normalized floating-point test data ranging from 0 to 1. Thus the normalized value of n_0 is 0.5748904732. A spreadsheet program is a convenient and reasonably universal way to generate these data. The frequency data may be converted to phase data by assuming an averaging time of 1, yielding a set of 1001 phase data points. Similarly, frequency offset and/or drift terms may be added to the data. These conversions can also be done by a spreadsheet program.

The values of this data set will be uniformly distributed between 0 and 1. While a data set with a normal (Gaussian) distribution would be more realistic, and could be produced by summing a number of independent uniformly distributed data sets, or by the Box-Muller method [5], this simpler data set is adequate for software validation.

Table 31. 1000-point frequency data set.

Averaging factor	1	10	100
# Data Points	1000	100	10
Maximum	9.957453e-01	7.003371e-01	5.489368e-01
Minimum	1.371760e-03	2.545924e-01	4.533354e-01
Average[a]	4.897745e-01	4.897745e-01	4.897745e-01
Median	4.798849e-01	5.047888e-01	4.807261e-01
Linear Slope[b,c]	6.490910e-06	5.979804e-05	1.056376e-03
Intercept[c]	4.865258e-01	4.867547e-01	4.839644e-01
Bisection Slope[b]	−6.104214e-06	−6.104214e-05	−6.104214e-04
1st Diff Slope[b]	1.517561e-04	9.648320e-04	1.011791e-03
Log Fit[b,d], a=	5.577220e-03	5.248477e-03	7.138988e-03
y(t)=a·ln (bt+1)+c, b=	9.737500e-01	4.594973e+00	1.420429e+02
y'(t)=ab/(bt+1), c=	4.570469e-01	4.631172e-01	4.442759e-01
Slope at end	5.571498e-06	5.237080e-05	7.133666e-04
Standard Dev[e]	2.884664e-01	9.296352e-02	3.206656e-02
Normal Allan Dev[f]	2.922319e-01	9.965736e-02	3.897804e-02
Overlap Allan Dev[h]	2.922319e-01	9.159953e-02	3.241343e-02
Mod Allan Dev[g,h]	2.922319e-01	6.172376e-02	2.170921e-02
Time Deviation[h]	1.687202e-01	3.563623e-01	1.253382e-00
Hadamard Deviation	2.943883e-01	1.052754e-01	3.910861e-02
Overlap Had Dev	2.943883e-01	9.581083e-02	3.237638e-02
Hadamard Total Dev	2.943883e-01	9.614787e-02	3.058103e-02
Total Deviation	2.922319e-01	9.134743e-02	3.406530e-02
Modified Total Dev	2.418528e-01	6.499161e-02	2.287774e-02
Time Total Deviation	1.396338e-01	3.752293e-01	1.320847e-00

Notes:
[a] Expected value = 0.5.
[b] All slopes are per interval.
[c] Least squares linear fit.
[d] Exact results will depend on iterative algorithm used. Data not suited to log fit.
[e] Sample (not population) standard deviation. Expected value = $1/\sqrt{12}$ = 0.2886751.
[f] Expected value equal to standard deviation for white FM noise.
[g] Equal to normal Allan deviation for averaging factor = 1.
[h] Calculated with listed averaging factors from the basic $\tau = 1$ data set.

Table 32. Error bars for n=1000 point =1 data set with avg. factor=10.

Allan Dev Type #Pts	Sigma Value Noise Type & Ratio	Conf Interval Remarks	χ^2 for 95% CF
Normal 99	9.965736e-02 W FM[a], B1=0.870	CI= 8.713870e-03[b,c] $\pm 1\sigma$ CI[f].	
Overlapping 981	9.159953e-02 W FM # χ^2 df = 146.177	Max $\sigma_y(\tau)$=1.014923e-01[g] Max $\sigma_y(\tau)$= 1.035201e-01[h] Min $\sigma_y(\tau)$= 8.223942e-02[h]	119.07 114.45 181.34
Modified[d] 972	6.172376e-02 W FM[e], R(n)=0.384 # χ^2 df = 94.620	Max $\sigma_y(\tau)$=7.044412e-02[g] Max $\sigma_y(\tau)$= 7.224944e-02[h] Min $\sigma_y(\tau)$= 5.419961e-02[h]	72.64 69.06 122.71

Notes:
[a] Theoretical B1=1.000 for W FM noise and 0.667 for F and W PM noise.
[b] Simple, noise-independent CI estimate $=\sigma_y(\tau)/\sqrt{N}$=1.001594e-02.
[c] This CI includes $\kappa(\alpha)$ factor that depends on noise type:

Noise	α	$\kappa(\alpha)$
W PM	2	0.99
F PM	1	0.99
W FM	0	0.87
F FM	-1	0.77
RW FM	-2	0.75

[d] BW factor $2\pi f_h \tau_0$ = 10. Applies only to F PM noise.
[e] Theoretical R(n) for W FM noise = 0.500 and 0.262 for F PM noise.
[f] Double-sided 68 % confidence interval: p = 0.158 and 0.842.
[g] Single-sided 95 % confidence interval: p = 0.950.
[h] Double-sided 95 % confidence interval: p = 0.025 and 0.975.

12.5. IEEE Standard 1139-1999

IEEE Standard 1139-1999, IEEE Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology – Random Instabilities contains several examples of stability calculations. Annex C.2 contains an example of an Allan deviation calculation, Annex C.3 has an example of a modified Allan deviation calculation, Annex C.4 has an example of a total deviation calculation, and Annex D contains examples of confidence interval calculations.

References for Software

1. W.J. Riley, "A Test Suite for the Calculation of Time Domain Frequency Stability," *Proc. 1995 IEEE Freq. Contrl. Symp.*, pp. 360-366, June 1995.
2. W.J. Riley, "Addendum to a Test Suite for the Calculation of Time Domain Frequency Stability," *Proc. 1996 IEEE Freq. Contrl. Symp.*, pp. 880-882, June 1996.
3. Stable32 Software for Frequency Stability Analysis, Hamilton Technical Services, Beaufort, SC 29907 USA, www.wriley.com.
4. "IEEE Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology - Random Instabilities," *IEEE Std 1139-1999*, July 1999.
5. W.H. Press, B.P. Flannery, S.A. Teukolsky and W.T. Vetterling, <u>Numerical Recipes in C</u>, ISBN 0-521-35465-X, Cambridge Univ. Press, Cambridge, U.K., 1988.

13   Glossary

The following terms are used in the field of frequency stability analysis

Aging	The change in frequency with time due to internal effects within the device.
Allan variance	The two-sample variance $\sigma^2_y(\tau)$ commonly used to measure frequency stability.
Averaging	The process of combining phase or frequency data into samples at a longer averaging time.
Averaging time	See Tau.
BW	Bandwidth, hertz.
Confidence limit	The uncertainty associated with a measurement. Often a 68 % confidence level or error bar.
Drift	The change in frequency with time due to all effects (including aging and environmental sensitivity).
Frequency data	A set of fractional frequency values, y[i], where i denotes equally-spaced time samples.
Hadamard variance	A three-sample variance, HVAR, that is similar to the two-sample Allan variance. It uses the second differences of the fractional frequencies, and is unaffected by linear frequency drift.
£(f)	$£(f) = 10 \cdot \log[\frac{1}{2} \cdot S_\phi(f)]$, the ratio of the SSB phase noise power in a 1 Hz BW to the total carrier power, dBc/Hz. Valid when noise power is much smaller than the carrier power.
MJD	The Modified Julian Date is based on the astronomical Julian Date, the number of days since noon on January 1, 4713 BC. The MJD is the Julian Date – 2 4000 000.5.
Modified sigma	A modified version of the Allan or total variance that uses phase averaging to distinguish between white and flicker PM noise processes.
MTIE	The maximum time interval error of a clock.
Normalize	To remove the average value from phase or frequency data.
Phase data	A set of time deviates, x[i] with units of seconds, where i denotes equally-spaced time samples. Called "phase" data to distinguish them from the independent time variable.
Phase noise	The spectral density of the phase deviations.
Sampling time	See Tau.
Sigma	The square root or deviation of a variance, often the two-sample or Allan deviation, $\sigma_y(\tau)$.
Slope	The change in frequency per tau interval.
SSB	Single sideband.
$S_f(f)$	The one-sided spectral density of the phase deviations, rad²/Hz.

$S_x(f)$	The one-sided spectral density of the time deviations, sec²/Hz.
$S_y(f)$	The one-sided spectral density of the fractional frequency deviations, 1/Hz.
Tau	The interval between phase measurements or the averaging time used for a frequency measurement.
Total	A variance using an extended data set that provides better confidence at long averaging times.
Thêo1	Theoretical variance #1, a variance providing stability data out to 75 % of the record length.
ThêoBR	A biased-removed version of Thêo1.
ThêoH	A hybrid combination of ThêoBR and the Allan variance.
TIE	The time interval error of a clock. Can be expressed as the rms time interval error TIE rms or the maximum time interval error MTIE.
Total variance	A two-sample variance similar to the Allan variance with improved confidence at large averaging factors.
x(t)	The instantaneous time deviation from a nominal time, $x(t) = \phi(t)/2\pi\nu_0$, s, where ν_0 is the nominal frequency, hertz. This dependent time variable is often called "phase" to distinguish it from the independent time variable t.
y(t)	The instantaneous fractional frequency deviation from a nominal frequency, $y(t) = [\nu(t)-\nu_0]/\nu_0 = x'(t)$, where ν_0 is the nominal frequency.

References for Glossary

1. "IEEE Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology - Random Instabilities," *IEEE Std 1139-1999*, July 1999.
2. Glossary of Time and Frequency Terms issued by Comite Consultatif International de Radio Communication – International Telecommunications (CCITT) Union, Geneva, Switzerland.

14   Bibliography

- **Notes**
  A. These references are arranged by topic and sorted by date.
  B. All of the FCS and UFFC references are available on-line for IEEE UFFC members.
  C. Only the 1990 and later PTTI proceedings are now available on-line.
  D. None of the EFTF or older IEEE I&M proceedings are available on-line.
  E. NIST papers are available on-line at http://tf.nist.gov/general/publications.htm.
  F. Hamilton Technical Services papers available on-line at http://www.wriley.com.

- **General**
  1. Special Issue on Frequency Stability, *Proc. IEEE*, Vol. 54, Feb. 1966.
  2. *Proc. IEEE*, Vol. 55, June 1967.
  3. J.A. Barnes, "Atomic Timekeeping and the Statistics of Precision Signal Generators," *IEEE Proceedings*, vol. 54, pp. 207-219, 1966.
  4. J.A. Barnes, et. al., "Characterization of Frequency Stability," *IEEE Trans. Instrum. Meas.*, Vol. IM-20, No. 2, pp. 105-120, May 1971.
  5. B.E. Blair (Editor), "Time and Frequency: Theory and Fundamentals," *NBS Monograph 140*, U.S. Department of Commerce, National Bureau of Standards, May 1974.
  6. G. Winkler, "A Brief Review of Frequency Stability Measures," *Proc. 8th PTTI Meeting*, pp. 489-527, Dec. 1976.
  7. J. Rutman, "Oscillator Specifications: A Review of Classical and New Ideas," *Proc. 31th Annu. Symp. on Freq. Contrl.*, pp.291-310, June 1977.
  8. J. Vanier and M. Tetu, "Time Domain Measurement of Frequency Stability," *Proc. 10th PTTI Meeting*, pp. 247-291, Nov. 1978.
  9. P.Lesage and C. Audoin, "Characterization and Measurement of Time and Frequency Stability", *Radio Science*, Vol. 14, No. 4, pp. 521-539, 1979.
  10. D.A. Howe, D.W. Allan and J.A. Barnes, "Properties of Signal Sources and Measurement Methods," *Proc. 35th Annu. Symp. on Freq. Contrl.*, pp. 1-47, May 1981. Also available on line at NIST web site.
  11. V.F. Kroupa (Editor), <u>Frequency Stability: Fundamentals and Measurement</u>, IEEE Press, Institute of Electrical and Electronic Engineers, New York, 1983, ISBN 0-87942-171-1.
  12. S.R. Stein, "Frequency and Time - Their Measurement and Characterization," Chapter 12, pp.191-416, <u>Precision Frequency Control</u>, Vol. 2, Edited by E.A. Gerber and A. Ballato, Academic Press, New York, 1985, ISBN 0-12-280602-6.
  13. *Proc. IEEE*, Vol. 74, Jan. 1986.
  14. C.A. Greenhall, "Frequency Stability Review," *Telecommunications and Data Acquisition Progress Report 42-88*, Oct-Dec 1986, Jet Propulsion Laboratory, Pasadena, CA, Feb. 1987.
  15. D.W. Allan, "Time and Frequency (Time-Domain) Characterization, Estimation, and Prediction of Precision Clocks and Oscillators", *IEEE Trans. Ultrasonics, Ferroelectrics and Freq. Contrl.*, Vol. UFFC-34, No. 6, pp. 647-654, Nov. 1987.
  16. D.W. Allan, et al., "Standard Terminology for Fundamental Frequency and Time Metrology," *Proc. 42nd Annu. Freq. Control Symp.*, pp. 419-425, June, 1988.
  17. D.B Sullivan, D.W Allan, D.A. Howe and F.L.Walls (Editors), "Characterization of Clocks and Oscillators," *NIST Technical Note 1337*, U.S. Department of Commerce, National Institute of Standards and Technology, March 1990.
  18. D.W. Allan, "Characterization of Precision Clocks and Oscillators," Class Notes, NRL Workshop, Nov. 1990.
  19. D.W. Allan, "Time and Frequency Metrology: Current Status and Future Considerations," *Proc. 5th European Freq. and Time Forum*, pp. 1-9, March 1991.
  20. Special Issue on Time and Frequency, *Proc. IEEE*, Vol. 79, July 1991.
  21. J. Rutman and F. L. Walls, "Characterization of Frequency Stability in Precision Frequency Sources," *Proc. IEEE*, Vol. 79, July 1991.

22. J.A. Barnes, "The Analysis of Frequency and Time Data," Austron, Inc., Dec. 1991.
23. C.A. Greenhall, "The Generalized Autocovariance: A Tool for Clock Statistics," *Telecommunications and Mission Operations Progress Report*, Vol. 42-137, Jet Propulsion Laboratory, Pasadena, CA, May 1999.
24. C. Audoin and B. Guinot, <u>The Measurement of Time</u>, Cambridge University Press, 2001
25. G.E.P. Box and G.M. Jenkins, <u>Time Series Analysis: Forecasting and Control</u>, San Francisco: Holden-Day, 1970.
26. W.J. Riley, "The Basics of Frequency Stability Analysis," Hamilton Technical Services.
27. W.J. Riley, "The Calculation of Time Domain Frequency Stability," Hamilton Technical Services.
28. W.J. Riley, "Techniques for Frequency Stability Analysis," Tutorial at the 2003 Intl. Freq. Cont. Symp., May 2003.

- **Standards and Specifications**

29. "Characterization of Frequency and Phase Noise," *Report 580*, International Radio Consultative Committee (C.C.I.R.), pp. 142-150, 1986.
30. MIL-PRF-55310, Oscillators, Crystal, General Specification For.
31. R.L. Sydnor (Editor), "The Selection and Use of Precise Frequency Systems," *ITU-R Handbook*, 1995.
32. *Guide to the Expression of Uncertainty in Measurement*, International Standards Organization, 1995, ISBN 92-67-10188-9.
33. "IEEE Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology - Random Instabilities," *IEEE Std 1139-1999*, July 1999.

- **Classic (Pre-Allan variance)**

34. Proc. IEEE-NASA Symp. on the Definition and Measurement of Short-Term Frequency Stability, *NASA SP-80*, Nov. 1964.

- **Allan Variance**

35. D.W. Allan, "Allan Variance," Allan's TIME, http://www.allanstime.com.
36. D.W. Allan, "The Statistics of Atomic Frequency Standards," *Proc. IEEE*, Vol. 54, No. 2, pp. 221-230, Feb. 1966.
37. "Characterization of Frequency Stability," *NBS Technical Note 394*, U.S Department of Commerce, National Bureau of Standards, Oct. 1970.
38. J.A. Barnes, et al, "Characterization of Frequency Stability," *IEEE Trans. Instrum. Meas.*, Vol. IM-20, No. 2, pp. 105-120, May 1971.
39. J.A. Barnes, "Variances Based on Data with Dead Time Between the Measurements," *NIST Technical Note 1318*, U.S. Department of Commerce, National Institute of Standards and Technology, 1990.
40. C.A. Greenhall, "Does Allan Variance Determine the Spectrum?" *Proc. 1997 Intl. Freq. Cont. Symp.*, pp. 358-365, June 1997.
41. C.A. Greenhall, "Spectral Ambiguity of Allan Variance," *IEEE Trans. Instrum. Meas.*, Vol. IM-47, No. 3, pp. 623-627, June 1998.
42. D.A. Howe, "Interpreting Oscillatory Frequency Stability Plots," *Proc. 2000 IEEE Freq. Contrl. Symp.*, pp. 725-732, May 2002

- **Modified Allan Variance**

43. D.W. Allan and J.A. Barnes, "A Modified Allan Variance with Increased Oscillator Characterization Ability," *Proc. 35th Annu. Symp. on Freq. Contrl.*, pp. 470-474, May 1981.

44. P. Lesage and T. Ayi, "Characterization of Frequency Stability: Analysis of the Modified Allan Variance and Properties of Its Estimate," *IEEE Trans. Instrum. Meas.*, Vol. IM-33, No. 4, pp. 332-336, Dec. 1984.

45. C.A. Greenhall, "Estimating the Modified Allan Variance," *Proc. 1995 IEEE Freq. Contrl. Symp.*, pp. 346-353, May 1995.

46. C.A. Greenhall, "The Third-Difference Approach to Modified Allan Variance," *IEEE Trans. Instrum. Meas.*, Vol. IM-46, No. 3, pp. 696-703, June 1997.

47. C.A. Greenhall, "A Shortcut for Computing the Modified Allan Variance," *Proc. 1992 IEEE Freq. Contrl. Symp.*, pp. 262-264, May 1992.

- **Time Variance and Time Error Prediction**

48. D.W. Allan and H. Hellwig, "Time Deviation and Time Prediction Error for Clock Specification, Characterization, and Application," *Proceedings of the Position Location and Navigation Symposium (PLANS)*, 29-36, 1978.

49. D.W. Allan, D.D. Davis, J. Levine, M.A. Weiss, N. Hironaka, and D. Okayama, "New Inexpensive Frequency Calibration Service From NIST," *Proc. 44th Annu. Symp. on Freq. Contrl.*, pp. 107-116, June 1990.

50. D.W. Allan, M.A. Weiss and J.L. Jespersen, "A Frequency-Domain View of Time-Domain Characterization of Clocks and Time and Frequency Distribution Systems," *Proc. 45th Annu. Symp. on Freq. Contrl.*, pp. 667-678, May 1991.

- **Hadamard Variance**

51. Jacques Saloman Hadamard (1865-1963), French mathematician.

52. W.K. Pratt, J. Kane and H.C. Andrews, "Hadamard Transform Image Coding," *Proc. IEEE*, Vol. 57, No. 1, pp.38-67, Jan. 1969.

53. R.A. Baugh, "Frequency Modulation Analysis with the Hadamard Variance," *Proc. 25th Annu. Symp. on Freq. Contrl.*, pp. 222-225, April 1971.

54. B. Picinbono, "Processus a Accroissements Stationnaires," *Ann. des telecom*, Tome 30, No. 7-8, pp. 211-212, July-Aug, 1975.

55. K. Wan, E. Visr and J. Roberts, "Extended Variances and Autoregressive Moving Average Algorithm for the Measurement and Synthesis of Oscillator Phase Noise," *43rd Annu. Symp. on Freq. Contrl.*, pp.331-335, June 1989.

56. S.T. Hutsell, "Relating the Hadamard Variance to MCS Kalman Filter Clock Estimation," *Proc. 27th PTTI Meeting*, pp. 291-302, Dec. 1995.

57. S.T. Hutsell, "Operational Use of the Hadamard Variance in GPS," *Proc. 28th PTTI Meeting*, pp. 201-213, Dec. 1996.

58. D.N. Matsakis and F.J. Josties, "Pulsar-Appropriate Clock Statistics," *Proc. 28th PTTI Meeting*, pp. 225-236, Dec. 1996.

59. W.J. Riley, "The Hadamard Variance," Hamilton Technical Services, 2006.

60. D. Howe, et al., "A Total Estimator of the Hadamard Function Used For GPS Operations," *Proc. 32nd PTTI Meeting*, Nov. 2000.

61. D.A. Howe, R.L. Beard, C.A. Greenhall, F. Vernotte, W.J. Riley and T.K. Peppler, "Enhancements to GPS Operations and Clock Evaluations Using a "Total" Hadamard Deviation," *IEEE Trans. Ultrasonics, Ferrelectrics and Freq. Contrl.*, Vol. 52, No. 8, pp. 1253-1261, Aug. 2005.

- **Modified Hadamard Variance**

62. S. Bregni and L. Jmoda, "Improved Estimation of the Hurst Parameter of Long-Range Dependent Traffic Using the Modified Hadamard Variance," *Proceedings of the IEEE ICC*, June 2006.

63. C.A. Greenhall and W.J. Riley, "Uncertainty of Stability Variances Based on Finite Differences," *Proceedings of the 35th Annual Precise Time and Time Interval (PTTI) Systems and Applications Meeting*, December 2003.

- **Total Variance**

64. D.A. Howe, "An Extension of the Allan Variance with Increased Confidence at Long Term," *Proc. 1995 IEEE Int. Freq. Cont. Symp.*, pp. 321-329, June 1995.
65. D.A. Howe and K.J. Lainson, "Simulation Study Using a New Type of Sample Variance," *Proc. 1995 PTTI Meeting*, pp. 279-290, Dec. 1995.
66. D.A. Howe and K.J. Lainson, "Effect of Drift on TOTALDEV," *Proc. 1996 Intl. Freq. Cont. Symp.* , pp. 883-889, June 1996
67. D.A. Howe, "Methods of Improving the Estimation of Long-term Frequency Variance," *Proc. 1997 European Frequency and Time Forum*, pp. 91-99, March 1997.
68. D.A. Howe & C.A. Greenhall, "Total Variance: A Progress Report on a New Frequency Stability Characterization," pp. 39-48, *Proc. 1997 PTTI Meeting*, Dec. 1997.
69. D.B. Percival and D.A. Howe, "Total Variance as an Exact Analysis of the Sample Variance," *Proc. 1997 PTTI Meeting*, Dec. 1997.
70. C.A. Greenhall, D.A. Howe and D.B Percival, "Total Variance, an Estimator of Long-Term Frequency Stability," *IEEE Trans. Ultrasonics, Ferroelectrics and Freq. Contrl.*, Vol. UFFC-46, No. 5, pp. 1183-1191, Sept. 1999.
71. D.A. Howe, "Total Variance Explained," *Proc. 1999 Joint Meeting of the European Freq. and Time Forum and the IEEE Freq. Contrl. Symp.*, pp. 1093-1099, April 1999.
72. D.A. Howe, "The Total Deviation Approach to Long-Term Characterization of Frequency Stability," *IEEE Trans. Ultrasonics, Ferroelectrics and Freq. Contrl.*, Vol. UFFC-47, No. 5, pp. 1102-1110, Sept. 2000.
73. D.A. Howe and T.K. Peppler, "Definitions of Total Estimators of Common Time-Domain Variances," *Proc. 2001 Intl. Freq. Cont. Symp.* , pp. 127-132, June 2001.

- **Modified Total Variance**

74. D.A. Howe and F. Vernotte, "Generalization of the Total Variance Approach to the Modified Allan Variance," *Proc. 31st PTTI Meeting*, pp.267-276, Dec. 1999.

- **Time Total Variance**

75. M.A. Weiss and D.A. Howe, "Total TDEV," *Proc. 1998 IEEE Int. Freq. Cont. Symp.*, pp. 192-198, June 1998. Note: An old definition of TDEV.
76. D.A. Howe and T.K. Peppler, "Definitions of Total Estimators of Common Time-Domain Variances," *Proc. 2001 Intl. Freq. Cont. Symp.* , pp. 127-132, June 2001.

- **Thêo1, ThêoBR, and ThêoH**

77. D.A. Howe and T.K. Peppler, "Very Long-Term Frequency Stability: Estimation Using a Special-Purpose Statistic," *Proceedings of the 2003 IEEE International Frequency Control Symposium*, pp. 233-238, May 2003.
78. D.A. Howe and T.N. Tasset, "Thêo1: Characterization of Very Long-Term Frequency Stability," *Proc. 2004 EFTF* .
79. D.A. Howe, "ThêoH: A Hybrid, High-Confidence Statistic that Improves on the Allan Deviation," *Metrologia* **43** (2006), S322-S331.
80. D.A. Howe, J. McGee-Taylor and T. Tasset, "ThêoH Bias-Removal Method," *Proc. 2006 IEEE Freq. Contrl. Symp.*, pp. 788-792, June 2006.

81. J. McGee and D.A. Howe, "ThêoH and Allan Deviation as Power-Law Noise Estimators," *IEEE Trans. Ultrasonics, Ferroelectrics and Freq. Contrl.*, Feb. 2007.
82. J. McGee and D.A. Howe, "Fast TheoBR: A method for long data set stability analysis," *IEEE Trans. Ultrasonics, Ferroelectrics and Freq. Contrl.*, to be published, 2008.

- **MTIE**

83. P. Travella and D. Meo, "The Range Covered by a Clock Error in the Case of White FM," *Proc. 30th Annu. PTTI Meeting*, pp. 49-60, Dec. 1998.
84. P. Travella, A. Dodone and S. Leschiutta, "The Range Covered by a Random Process and the New Definition of MTIE," *Proc. 28th Annu. PTTI Meeting*, pp. 119-123, Dec. 1996.
85. S. Bregni, "Clock Stability Characterization and Measurement in Tele-communications," *IEEE Trans. Instrum. Meas.*, Vol. IM-46, No. 6, pp. 1284-1294, Dec. 1997.
86. S. Bregni, "Measurement of Maximum Time Interval Error for Telecommunications Clock Stability Characterization," *IEEE Trans. Instrum. Meas.*, Vol. IM-45, No. 5, pp. 900-906, Oct. 1996.
87. G. Zampetti, "Synopsis of Timing Measurement Techniques Used in Telecommunucations," *Proc. 24th PTTI Meeting*, pp. 313-326, Dec. 1992.
88. M.J. Ivens, "Simulating the Wander Accumulation in a SDH Synchronisation Network," Master's Thesis, University College, London, UK, November 1997.
89. S. Bregni and S. Maccabruni, "Fast Computation of Maximum Time Interval Error by Binary Decomposition," IEEE Trans. I&M, Vol. 49, No. 6, pp. 1240-1244, Dec. 2000.

- **Multi-Variance Analysis**

90. F. Vernotte, E. Lantz, "Time Stability: An Improvement of the Multi-Variance Method for the Oscillator Noise Analysis," *Proc. 6th European Frequency and Time Forum*, pp. 343-347, March 1992.
91. F. Vernotte, E. Lantz, J. Groslambert and J.J. Gagnepain, "A New Multi-Variance Method for the Oscillator Noise Analysis," *Proc. 46th Annu. Symp. on Freq. Contrl.*, pp. 284-289, May 1992.
92. F. Vernotte, E. Lantz, J. Groslambert and J.J. Gagnepain, "Oscillator Noise Analysis: Multivariance Measurement," *IEEE Trans. Instrum. Meas.*, Vol. IM-42, No. 2, pp. 342-350, April 1993.
93. F. Vernotte, E. Lantz, F. Meyer and F. Naraghi, "Simultaneneous Measurement of Drifts and Noise Coefficients of Oscillators: Application to the Analysis of the Time Stability of the Millisecond Pulsars," *Proc. 1997 European Frequency and Time Forum*, pp. 91-99, March 1997.
94. T. Walter, "A Multi-Variance Analysis in the Time Domain," *Proc. 24th PTTI Meeting*, pp. 413-424, Dec. 1992.

- **Dynamic Stability**

95. L. Galleani and P. Tavella, "The Characterization of Clock Behavior with the Dynamic AllanVariance," *Proc. 2003 Joint FCS/EFTF Meeting*, pp. 239-244.
96. L. Galleani and P. Tavella, "Tracking Nonstationarities in Clock Noises Using the Dynamic Allan Variance," *Proc. 2005 Joint FCS/PTTI Meeting*.

- **Confidence Intervals**

97. J.A. Barnes and D.W. Allan, "Variances Based on Data with Dead Time Between the Measurements," *NIST Technical Note 1318*, U.S. Department of Commerce, National Institute of Standards and Technology, 1990.
98. C.A. Greenhall, "Recipes for Degrees of Freedom of Frequency Stability Estimators," *IEEE Trans. Instrum. Meas.*, Vol. 40, No. 6, pp. 994-999, Dec. 1991.
99. M.A. Weiss and C. Hackman, "Confidence on the Three-Point Estimator of Frequency Drift," *Proc. 24th Annu. PTTI Meeting*, pp. 451-460, Dec. 1992.

100.   M.A.Weiss and C.A. Greenhall, "A Simple Algorithm for Approximating Confidence on the Modified Allan Variance and the Time Variance," *Proc. 28th Annu. PTTI Meeting*, pp. 215-224, Dec. 1996.
101.   D.A. Howe, "Methods of Improving the Estimation of Long-Term Frequency Variance," *Proc. 11th European Freq. and Time Forum*, pp. 91-99, March 1997.
102.   W.J. Riley, "Confidence Intervals and Bias Corrections for the Stable32 Variance Functions," Hamilton Technical Services, 2000.
103.   F. Vernotte and M. Vincent, "Estimation of the Uncertainty of a Mean Frequency Measurement," *Proc. 11th European Freq. and Time Forum*, pp. 553-556, March 1997.
104.   P. Lesage and C. Audoin, "Estimation of the Two-Sample Variance with a Limited Number of Data," pp. 311-318.
105.   P. Lesage and C. Audoin, "Characterization of Frequency Stability: Uncertainty due to the Finite Number of Measurements," *IEEE Trans. Instrum. Meas.*, Vol. 22, No. 2, pp.157-161, June 1973.
106.   K. Yoshimura, "Degrees of Freedom of the Estimate of the Two-Sample Variance in the Continuous Sampling Method," *IEEE Trans. Instrum. Meas.*, Vol. 38, No. 6, pp. 1044-1049, Dec. 1989.
107.   C.R. Ekstrom and P.A. Koppang, "Error Bars for Three-Cornered Hats," *IEEE Trans. Ultrasonics, Ferroelectrics and Freq. Contrl.*, Vol. 53, No. 5, pp. 876-879, May 2006.
108.   C. Greenhall and W. Riley, "Uncertainty of Stability Variances Based on Finite Differences," *Proc. 35th PTTI Meeting*, Dec. 2003.
109.   T.N. Tasset, D.A. Howe and D.B. Percival, "Thêo1 Confidence Intervals," *Proc. 2004 Joint FCS/UFFC Meeting*, pp. 725-728, Aug. 2004.

- **Drift Estimation and Removal**

110.   C.A. Greenhall, "Removal of Drift from Frequency Stability Measurements," *Telecommunications and Data Acquisition Progress Report 42-65*, July-Aug 1981, Jet Propulsion Laboratory, Pasadena, CA, pp. 127-132, 1981.
111.   J.A. Barnes, "The Measurement of Linear Frequency Drift in Oscillators," *Proc. 15th Annu. PTTI Meeting*, pp. 551-582, Dec. 1983.
112.   M.A. Weiss and C. Hackman, "Confidence on the Three-Point Estimator of Frequency Drift," *Proc. 24th Annu. PTTI Meeting*, pp. 451-460, Dec. 1992.
113.   M.A. Weiss, D.W. Allan and D.A. Howe, "Confidence on the Second Difference Estimation of Frequency Drift," *Proc. 1992 IEEE Freq. Contrl. Symp.*, pp. 300-305, June 1992.
114.   "Long Term Quartz Oscillator Aging - A New Approach," *The Standard*, Hewlett-Packard Co., pp. 4-5, Jan. 1994.
115.   L.A. Breakiron, "A Comparative Study of Clock Rate and Drift Estimation."
116.   G.Wei, "Estimations of Frequency and its Drift Rate," *IEEE Trans. Instrum. Meas.*, Vol. 46, No. 1, pp. 79-82, Feb. 1997.
117.   C.A. Greenhall, "A Frequency-Drift Estimator and Its Removal from Modified Allan Variance," *Proc. 1997 IEEE Freq. Contrl. Symp.*, pp. 428-432, June 1997.
118.   F. Vernotte and M. Vincent, "Estimation of the Measurement Uncertainty of Drift Coefficients Versus the Noise Levels," *Proc. 12th European Freq. and Time Forum*, pp. 222-227, March 1998.

- **Bias Functions**

119.   J.A. Barnes, "Tables of Bias Functions, B_1 and B_2, for Variances Based on Finite Samples of Processes with Power Law Spectral Densities," *NBS Technical Note 375*, Jan. 1969.
120.   J.A. Barnes, "The Analysis of Frequency and Time Data," Austron, Inc., Dec. 1991.
121.   W.J. Riley, "Confidence Intervals and Bias Corrections for the Stable32 Variance Functions," Hamilton Technical Services

- **Noise Identification**

122. J.A. Barnes, "The Generation and Recognition of Flicker Noise," *NBS Report 9284*, U.S Department of Commerce, National Bureau of Standards, June 1967.

123. J.A. Barnes, "Effective Stationarity and Power Law Spectral Densities," NBS Frequency-Time Seminar Preprint, Feb. 1968.

124. J.A. Barnes and D.W. Allan, "Recognition and Classification of LF Divergent Noise Processes," NBS Division 253 Class Notes, circa 1970.

125. J.A. Barnes, "Models for the Interpretation of Frequency Stability Measurements," *NBS Technical Note 683*, U.S Department of Commerce, National Bureau of Standards, Aug. 1976.

126. C.A. Greenhall and J.A. Barnes, "Large Sample Simulation of Flicker Noise," *Proc. 19th Annu. PTTI Meeting*, pp. 203-217, Dec. 1987 and *Proc. 24th Annu. PTTI Meeting*, p. 461, Dec. 1992.

127. D. Howe, R. Beard, C. Greenhall, F. Vernotte and W. Riley, "A Total Estimator of the Hadamard Function Used for GPS Operations," *Proc. 32nd PTTI Meeting*, pp. 255-268, Nov. 2000.

128. W.J. Riley and C.A. Greenhall, "Power Law Noise Identification Using the Lag 1 Autocorrelation," *Proc. 18th European Freq. and Time Forum*, April 2004.

129. W.J. Riley, "Confidence Intervals and Bias Corrections for the Stable32 Variance Functions," Hamilton Technical Services.

- **Simulation**

130. J.A Barnes and D.W. Allan, "A Statistical Model of Flicker Noise," *Proc. IEEE*, Vol. 54, No. 2, pp. 176-178, Feb. 1966.

131. J.A. Barnes, "The Generation and Recognition of Flicker Noise," *NBS Report 9284*, U.S Department of Commerce, National Bureau of Standards, June 1967.

132. N.J. Kasdin and T. Walter, "Discrete Simulation of Power Law Noise," *Proc. 1992 IEEE Freq. Contrl. Symp.*, pp. 274-283, May 1992.

133. T. Walter, "Characterizing Frequency Stability: A Continuous Power-Law Model with Discrete Sampling," *IEEE Trans. Instrum. Meas.*, Vol. 43, No. 1, pp. 69-79, Feb. 1994.

134. S.K. Park and K.W. Miller, "Random Number Generators: Good Ones are Hard to Find," *Comm. ACM*, Vol. 31, No. 10, pp. 1192-1201.

135. S. Bregni, L. Valtriani, F. Veghini, "Simulation of Clock Noise and AU-4 Pointer Action in SDH Equipment," *Proc. of IEEE GLOBECOM '95*, pp. 56-60, Nov. 1995.

136. S. Bregni, "Generation of Pseudo-Random Power-Law Noise Sequences by Spectral Shaping" in *Communications World*. Editor: N. Mastorakis. WSES Press, 2001.

- **Dead Time**

137. J.A. Barnes and D.W. Allan, "Variances Based on Data with Dead Time Between the Measurements," *NIST Technical Note 1318*, U.S. Department of Commerce, National Institute of Standards and Technology, 1990.

138. D.B Sullivan, D.W Allan, D.A. Howe and F.L. Walls (Editors), "Characterization of Clocks and Oscillators," *NIST Technical Note 1337*, U.S. Department of Commerce, National Institute of Standards and Technology, TN-296-335, March 1990.

139. J.A. Barnes, "The Analysis of Frequency and Time Data," Austron, Inc., Dec. 1991.

140. D.A. Howe and E.E. Hagn, "Limited Live-Time Measurements of Frequency Stability," *Proc. Joint FCS/EFTF*, pp. 1113-1116, April 1999.

141. W.J. Riley, "Gaps, Outliers, Dead Time, and Uneven Spacing in Frequency Stability Data," Hamilton Technical Services.

- **3-Cornered Hat**

142.	J.E. Gray and D.W. Allan, "A Method for Estimating the Frequency Stability of an Individual Oscillator," *Proc. 28th Freq. Contrl. Symp.*, pp. 243-246, May 1974.
143.	J. Groslambert, D. Fest, M. Oliver and J.J. Gagnepain, "Characterization of Frequency Fluctuations by Crosscorrelations and by Using Three or More Oscillators," *Proc..35th Freq. Contrl. Symp.*, pp. 458-462, May 1981.
144.	S.R. Stein, "Frequency and Time - Their Measurement and Characterization", Chapter 12, Section 12.1.9, *Separating the Variances of the Oscillator and the Reference*, pp. 216-217, <u>Precision Frequency Control</u> , Vol. 2, Edited by E.A. Gerber and A. Ballato, Academic Press, New York, 1985, ISBN 0-12-280602-6.
145.	P. Tavella and A. Premoli, ""Characterization of Frequency Standard Instability by Estimation of their Covariance Matrix," *Proc. 23rd PTTI Meeting*, pp. 265-276. Dec. 1991.
146.	P. Tavella and A. Premoli, "A Revisited Tree-Cornered Hat Method for Estimating Frequency Standard Instability," *IEEE Trans. Instrum. Meas.*, IM-42, pp. 7-13, Feb. 1993.
147.	C.R. Ekstrom and P.A. Koppang, "Error Bars for Three-Cornered Hats," *IEEE Trans. UFFC*, Vol. 53, No. 5, pp. 876-879, May 2006.
148.	W.J. Riley, "Application of the 3-Cornered Hat Method to the Analysis of Frequency Stability," Hamilton Technical Services.

- **Domain Conversions**

149.	A.R. Chi, "The Mechanics of Translation of Frequency Stability Measures Between Frequency and Time Domain Measurements," *Proc. 9th Annu. PTTI Meeting*, pp. 523-548, Dec.1977.
150.	J. Rutman, "Relations Between Spectral Purity and Frequency Stability," pp. 160-165.
151.	R. Burgoon and M.C. Fisher, "Conversion between Time and Frequency Domain of Intersection Points of Slopes of Various Noise Processes," *32th Annu. Symp. on Freq. Contrl.*, pp. 514-519, June 1978.
152.	W.F. Egan, "An Efficient Algorithm to Compute Allan Variance from Spectral Density," *IEEE Trans. Instrum. Meas.*, Vol. 37, No. 2, pp. 240-244, June 1988.
153.	F. Vernotte, J. Groslambert and J.J. Gagnepain, "Practical Calculation Methods of Spectral Density of Instantaneous Normalized Frequency Deviation from Instantaneous Time Error Samples," *Proc. 5th European Freq. and Time Forum*, pp. 449-455, March 1991.
154.	F. Thomson, S. Asmar and K. Oudrhiri, "Limitations on the Use of the Power-Law Form of Sy(f) to Compute Allan Variance," *IEEE Trans. Ultrasonics, Ferroelectrics and Freq. Contrl.*, Vol. 52, No. 9, pp. 1468-1472, Sept. 2005.
155.	W.J. Riley, "Stable32 Frequency Domain Functions," Hamilton Technical Services.

- **Robust Statistics**

156.	D.B. Percival, "Use of Robust Statistical Techniques in Time Scale Formation," Preliminary Report, U.S. Naval Observatory Contract No. N70092-82-M-0579, 1982.
157.	Gernot M.R. Winkler, "Introduction to Robust Statistics and Data Filtering," Tutorial at 1993 IEEE Freq. Contrl. Symp., Sessions 3D and 4D, June 1, 1993.
158.	V. Barnett and T. Lewis, *Outliers in Statistical Data*, 3[rd] Edition, John Wiley & Sons, Chichester, 1994, ISBN 0-471-93094-6.

- **Computation and Algorithms**

159.	W.H. Press, B.P. Flannery, S.A. Teukolsky and W.T. Vetterling, *Numerical Recipes in C*, Cambridge Univ. Press, Cambridge, U.K., 1988, pp.216-217.
160.	C.A. Greenhall, "A Shortcut for Computing the Modified Allan Variance," *Proc. 1992 IEEE Freq. Contrl. Symp.*, pp. 262-264, May 1992.
161.	W.J. Riley, "A Test Suite for the Calculation of Time Domain Frequency Stability," *Proc. 1995 IEEE Freq. Contrl. Symp.*, pp. 360-366, June 1995.

162.	W.J. Riley, "Addendum to a Test Suite for the Calculation of Time Domain Frequency Stability," *Proc. 1996 IEEE Freq. Contrl. Symp.*, pp. 880-882, June 1996.
163.	M. Kasznia, "Some Approach to Computation of ADEV, TDEV and MTIE," *Proc. 11th European Freq. and Time Forum*, pp. 544-548, March 1997.

● **Measurements**

164.	D.W. Allan, "Picosecond Time Difference Measurement System," *Proc. 29th Annu. Symp. on Freq. Contrl.*, pp. 404-411, May 1975.
165.	D.A. Howe, D.W. Allan and J.A. Barnes, "Properties of Signal Sources and Measurement Methods", *Proc. 35th Annu. Symp. on Freq. Contrl.*, May 1981, pp. 1-47.
166.	S.R. Stein, "Frequency and Time - Their Measurement and Characterization," Chapter 12, pp.191-416, *Precision Frequency Control*, Vol. 2, Edited by E.A. Gerber and A. Ballato, Academic Press, New York, 1985, ISBN 0-12-280602-6.
167.	S. Stein, D. Glaze, J. Levine, J. Gray, D. Hilliard, D. Howe and L Erb, "Performance of an Automated High Accuracy Phase Measurement System," *Proc. 36th Annu. Freq. Contrl. Symp.*, June 1982, pp. 314-320.
168.	S.R. Stein and G.A. Gifford, "Software for Two Automated Time Measurement Systems," *Proc. 38th Annu. Freq. Contrl.* Symp, June 1984, pp. 483-486.
169.	Data Sheet, Model 5110A Time Interval Analyzer, Symmetricom, Inc. (formerly Timing Solutions Corp.), Boulder, CO 80301 USA.
170.	Data Sheet, A7 Frequency & Phase Comparator Measurement System, Quartzlock (UK) Ltd., Gothic, Plymouth Road, Totnes, Devon, TQ9 5LH England.
171.	C.A. Greenhall, "Oscillator-Stability Analyzer Based on a Time-Tag Counter," *NASA Tech Briefs*, NPO-20749, May 2001, p. 48.
172.	C.A. Greenhall, "Frequency Stability Review," Telecommunications and Data Acquisition Progress Report 42-88, Oct-Dec 1986, Jet Propulsion Laboratory, Pasadena, CA, pp. 200-212, Feb. 1987.

Index

About the Author

Mr. Riley has worked in the area of frequency control his entire professional career. He is currently the Proprietor of Hamilton Technical Services, where he provides software and consulting services in that field, including the Stable program for the analysis of frequency stability. Bill collaborates with other experts in the time and frequency community to provide an up-to-date and practical tool for frequency stability analysis that has received wide acceptance within that community. From 1999 to 2004, he was Manager of Rubidium Technology at Symmetricom, Inc. (previously Datum), applying his extensive experience with atomic frequency standards to those products within that organization, including the early development of a chip-scale atomic clock (CSAC). From 1980-1998, Mr. Riley was the Engineering Manager of the Rubidium Department at EG&G (now PerkinElmer), where his major responsibility was to direct the design of rubidium frequency standards, including high-performance rubidium clocks for the GPS navigational satellite program. Other activities there included the development of several tactical military and commercial rubidium frequency standards. As a Principal Engineer at Harris Corporation, RF Communications Division in 1978-1979, he designed communications frequency synthesizers. From 1962-1978, as a Senior Development Engineer at GenRad, Inc. (previously General Radio), Mr. Riley was responsible for the design of electronic instruments, primarily in the area of frequency control. He has a 1962 BSEE degree from Cornell University and a 1966 MSEE degree from Northeastern University. Mr. Riley holds six patents in the area of frequency control, and has published a number of papers and tutorials in that field. He is a Fellow of the IEEE, and a member of Eta Kappa Nu, the IEEE UFFC Society, and served on the PTTI Advisory Board. He received the 2000 IEEE International Frequency Control Symposium I.I. Rabi Award for his work on atomic frequency standards and frequency stability analysis.

William J. Riley, Jr.
Proprietor
Hamilton Technical Services
650 Distant Island Drive
Beaufort, SC 29907
Phone: 843-525-6495
Fax: 843-525-0251
E-Mail: bill@wriley.com
Web: www.wriley.com

NIST Technical Publications

Periodical

Journal of Research of the National Institute of Standards and Technology–Reports NIST research and development in metrology and related fields of physical science, engineering, applied mathematics, statistics, biotechnology, and information technology. Papers cover a broad range of subjects, with major emphasis on measurement methodology and the basic technology underlying standardization. Also included from time to time are survey articles on topics closely related to the Institute's technical and scientific programs. Issued six times a year.

Nonperiodicals

Monographs–Major contributions to the technical literature on various subjects related to the Institute's scientific and technical activities.

Handbooks–Recommended codes of engineering and industrial practice (including safety codes) developed in cooperation with interested industries, professional organizations, and regulatory bodies.

Special Publications–Include proceedings of conferences sponsored by NIST, NIST annual reports, and other special publications appropriate to this grouping such as wall charts, pocket cards, and bibliographies.

National Standard Reference Data Series–Provides quantitative data on the physical and chemical properties of materials, compiled from the world's literature and critically evaluated. Developed under a worldwide program coordinated by NIST under the authority of the National Standard Data Act (Public Law 90-396). NOTE: The Journal of Physical and Chemical Reference Data (JPCRD) is published bimonthly for NIST by the American Institute of Physics (AlP). Subscription orders and renewals are available from AIP, P.O. Box 503284, St. Louis, MO 63150-3284.

Building Science Series–Disseminates technical information developed at the Institute on building materials, components, systems, and whole structures. The series presents research results, test methods, and performance criteria related to the structural and environmental functions and the durability and safety characteristics of building elements and systems.

Technical Notes–Studies or reports which are complete in themselves but restrictive in their treatment of a subject. Analogous to monographs but not so comprehensive in scope or definitive in treatment of the subject area. Often serve as a vehicle for final reports of work performed at NIST under the sponsorship of other government agencies.

Voluntary Product Standards–Developed under procedures published by the Department of Commerce in Part 10, Title 15, of the Code of Federal Regulations. The standards establish nationally recognized requirements for products, and provide all concerned interests with a basis for common understanding of the characteristics of the products. NIST administers this program in support of the efforts of private-sector standardizing organizations.

*Order the **following** NIST publications–FIPS and NISTIRs–from the National Technical Information Service, Springfield, VA 22161.*

Federal Information Processing Standards Publications (FIPS PUB)–Publications in this series collectively constitute the Federal Information Processing Standards Register. The Register serves as the official source of information in the Federal Government regarding standards issued by NIST pursuant to the Federal Property and Administrative Services Act of 1949 as amended, Public Law 89-306 (79 Stat. 1127), and as implemented by Executive Order 11717 (38 FR 12315, dated May 11,1973) and Part 6 of Title 15 CFR (Code of Federal Regulations).

NIST Interagency or Internal Reports (NISTIR)–The series includes interim or final reports on work performed by NIST for outside sponsors (both government and nongovernment). In general, initial distribution is handled by the sponsor; public distribution is handled by sales through the National Technical Information Service, Springfield, VA 22161, in hard copy, electronic media, or microfiche form. NISTIR's may also report results of NIST projects of transitory or limited interest, including those that will be published subsequently in more comprehensive form.

U.S. Department of Commerce
National Institute of Standards and Technology
325 Broadway
Boulder, CO 80305-3337

Official Business
Penalty for Private Use $300