Measurement of Ring Oscillator Noise & Analysis Using the Allan Variance Method

Rich Newell*
Microsemi System-on-Chip Products Group (formerly Actel Corp.)
Cryptarchi Conference
June 16th, 2011

*Special thanks to Mir Sayed Ali and Prasad Kuruganti for assistance with the experimental set-up

Motivation

- Ring oscillator phase noise is often used as a source of entropy for true random number generators
  - Å By capturing the random accumulated phase difference between two osc.

- The motivation of this analysis is to characterize the noise process, to be able to better answer questions like:
  - Å What is the optimal time to wait before sampling?
  - Å How does the variation in phase increase with time?
  - Å What mathematical models best describe the random process?
  - Å What physical processes are responsible for the oscillator behavior?

- Many authors in cryptology assume that ring oscillator
  ñflipping times [are] independent and identically distributedò[1]
  which would lead to a ñrandom walk in phaseòcharacteristic
  ï Is this assumption justified?

- [1]For example: ñOn the security of oscillator-based random number generators,òMathieu Baudet, et al

Scope

- Gather some data from various types of on-chip FPGA oscillators
- Demonstrate the use of the Allan Variance and related time-domain methods to characterize oscillator noise processes
  - Empirically identify dominant power-law slopes for the oscillators tested
- Proposing physical models for the types of oscillators tested is beyond the scope of this paper

Phase vs. Frequency

- Phase is the integral of frequency
  - Or, conversely, frequency is the derivative of phase

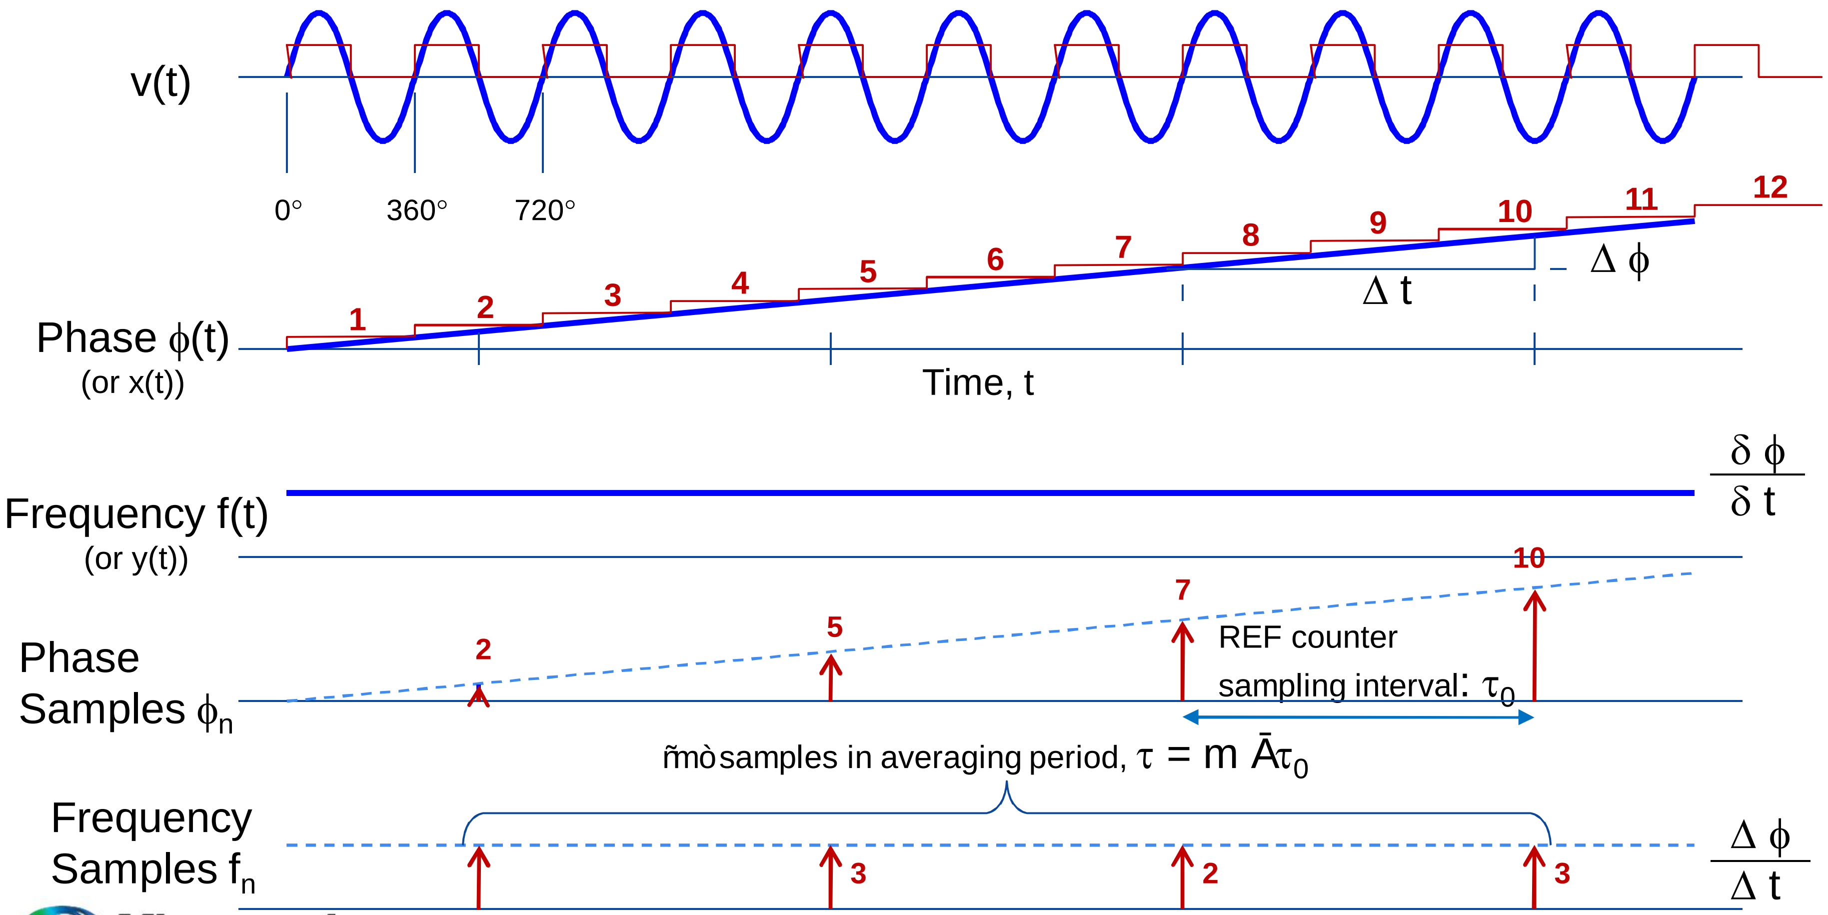

Frequency-domain Power-Law Slopes

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

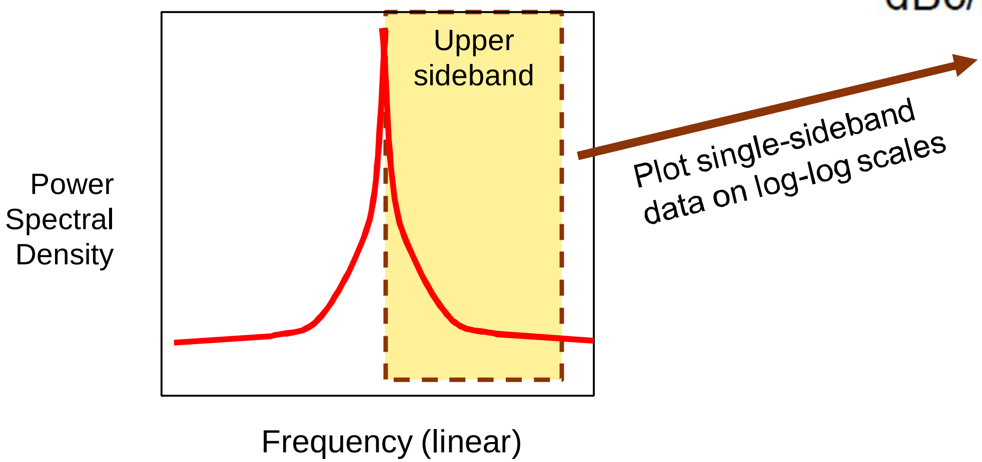

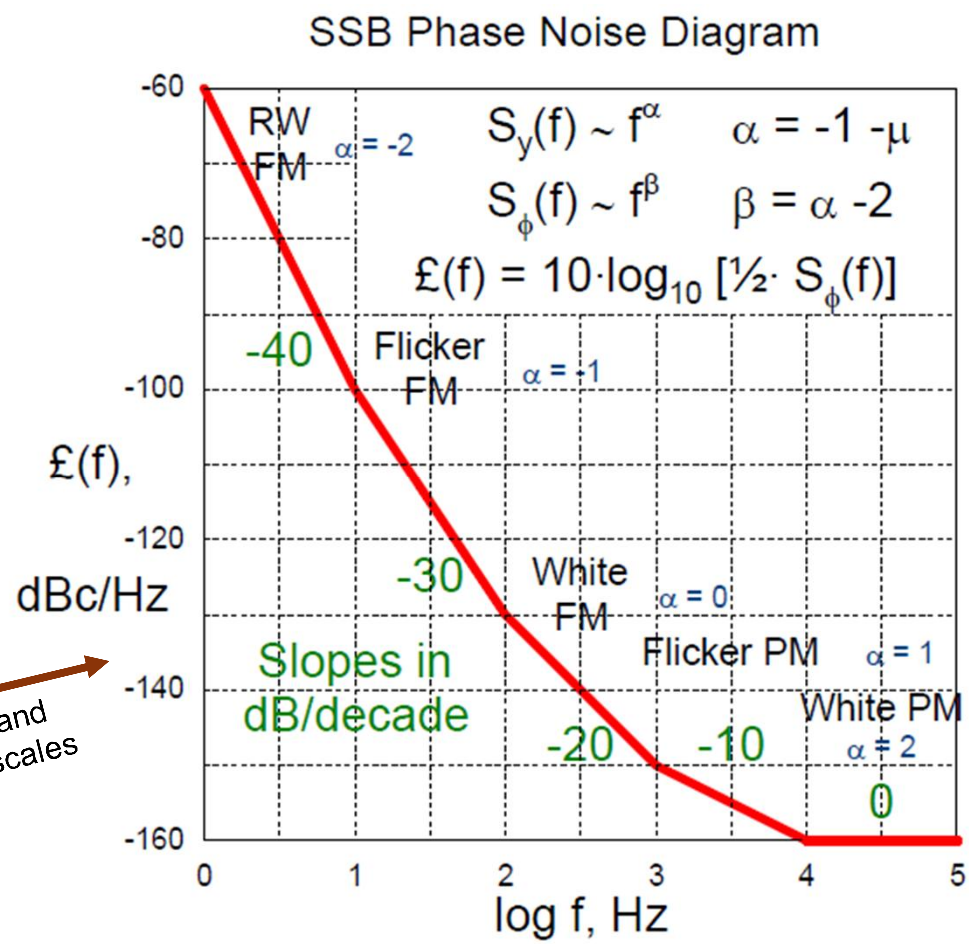

PM = Phase Modulation
FM = Frequency Modulation
RW = Random Walk
SSB = Single Side-Band

Problems with Frequency-Domain Analysis

- Needs expensive equipment
  - Usually requires a spectrum analyzer
  - Often needs low-noise PLL, mixers, etc. to heterodyne the signal to DC so that single-sideband measurements can be taken

- Not well suited for clocks
  - Hard to separate amplitude noise from phase noise (need to compare phase data in both sidebands)… we are uninterested in amplitude
  - Harmonics from square-waves can cause measurement issues
  - Spectrum analyzers not geared for large number of samples taken over a long time periods such as minutes, hours, days, or years
    - Most are designed for RF signals

Time-Domain Analysis “Equipment”*

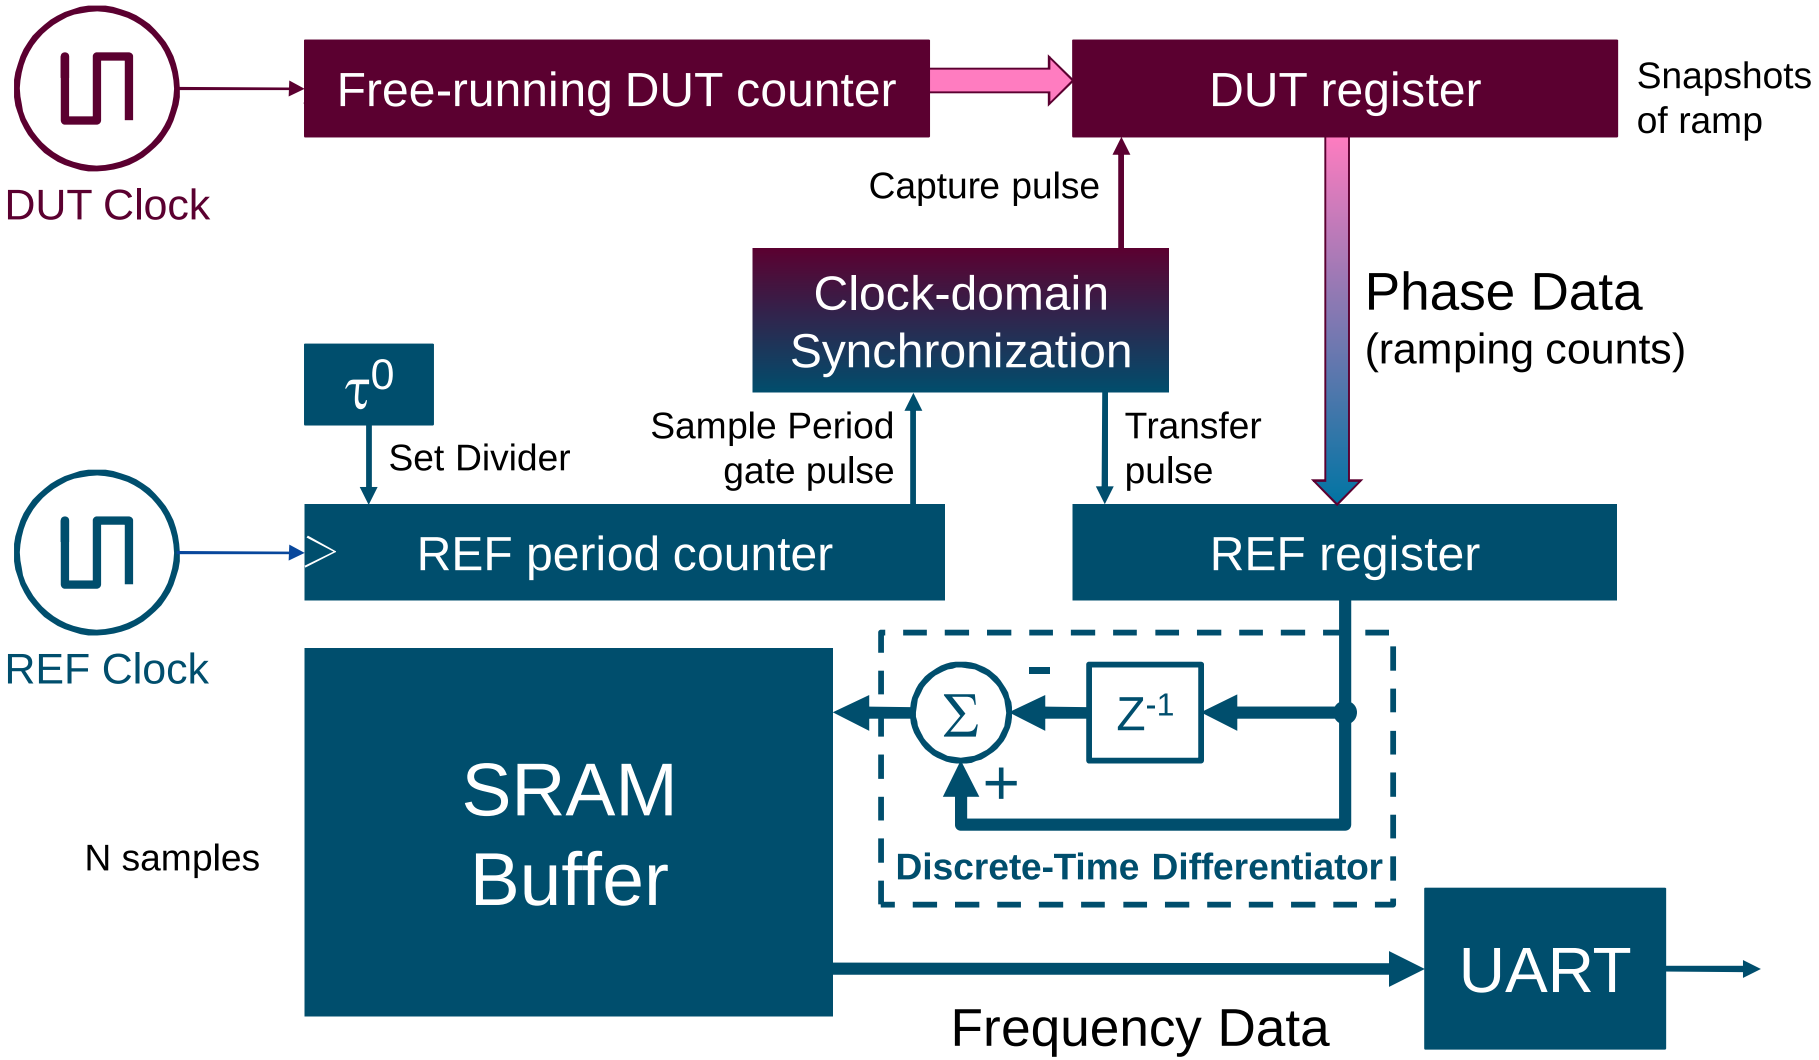

***All implemented inside FPGA**

Time-domain Measures of Frequency Stability

Frequency drift over time (for example, a random walk in frequency):

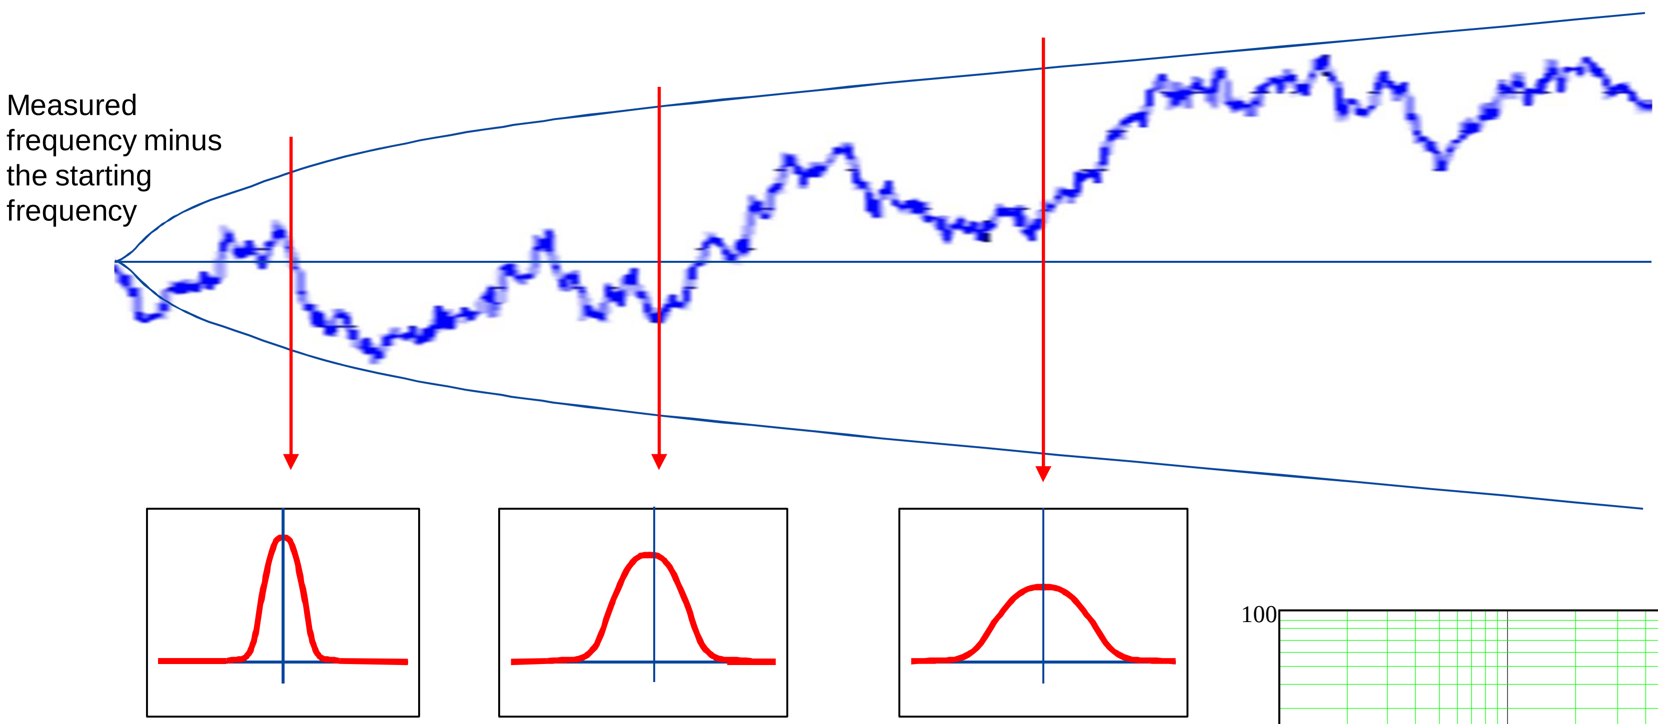

For a random walk, the standard deviation (σ) of an ensemble grows as the square-root of the averaging time (τ)

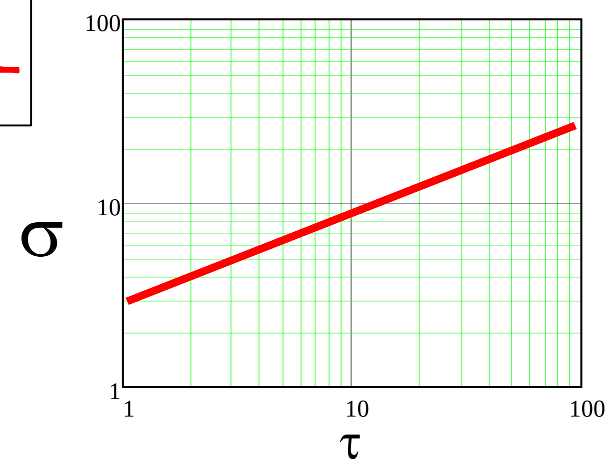

Time Domain Example Waveforms

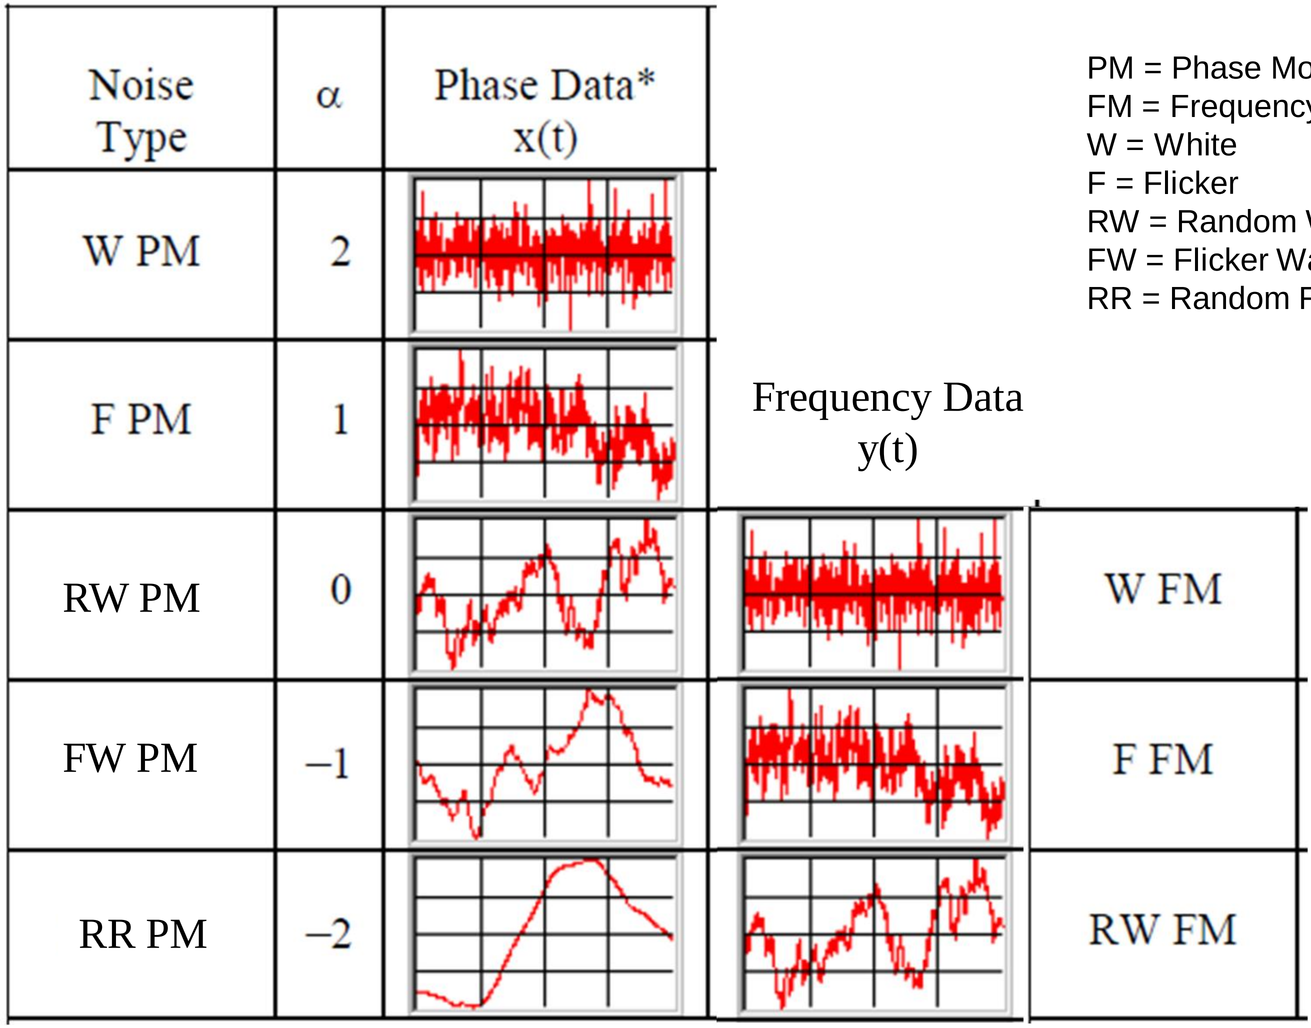

PM = Phase Modulation
FM = Frequency Modulation
W = White
F = Flicker
RW = Random Walk
FW = Flicker Walk
RR = Random Run

* Differencing the Phase Data generates Frequency data that looks similar to that shown two rows higher in the chart

Allan Variance

Standard Variance:

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

Allan Variance:
(in terms of fractional-frequency samples, $\bar{y}_i$, averaged over the period τ)

$$\sigma_y^2(\tau) = \frac{1}{2(M-1)}\sum_{i=1}^{M-1}[\overline{y}_{i+1} - \overline{y}_i]^2$$

First difference of frequency averages

N is total number of phase samples

M is N-1

Allan Variance:
(in terms of phase samples, x_i)

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

Second difference of phase samples

Calculated for multiple averaging periods: $\tau = m\tau_0$

Standard Variance: Squared deviation from meané divergent if over longer averaging intervals if the mean is ñwanderingò due to higher-order noise terms such as a random walk in frequency, or linear frequency drift

The Allan Variance: Based upon first difference (discrete-time derivative) of frequency samples instead, thus rejecting linear frequency drift This is the same as the second difference of phase samples

Time-Domain Power-Law Slopes: σ-τ log-log Plots

Allan Variance* and Modified Allan Variance* sigma-tau log-log Plots

τ = averaging time (x-axis) = m Āτ_0, the sampling time

σ = Standard Deviation (y-axis)

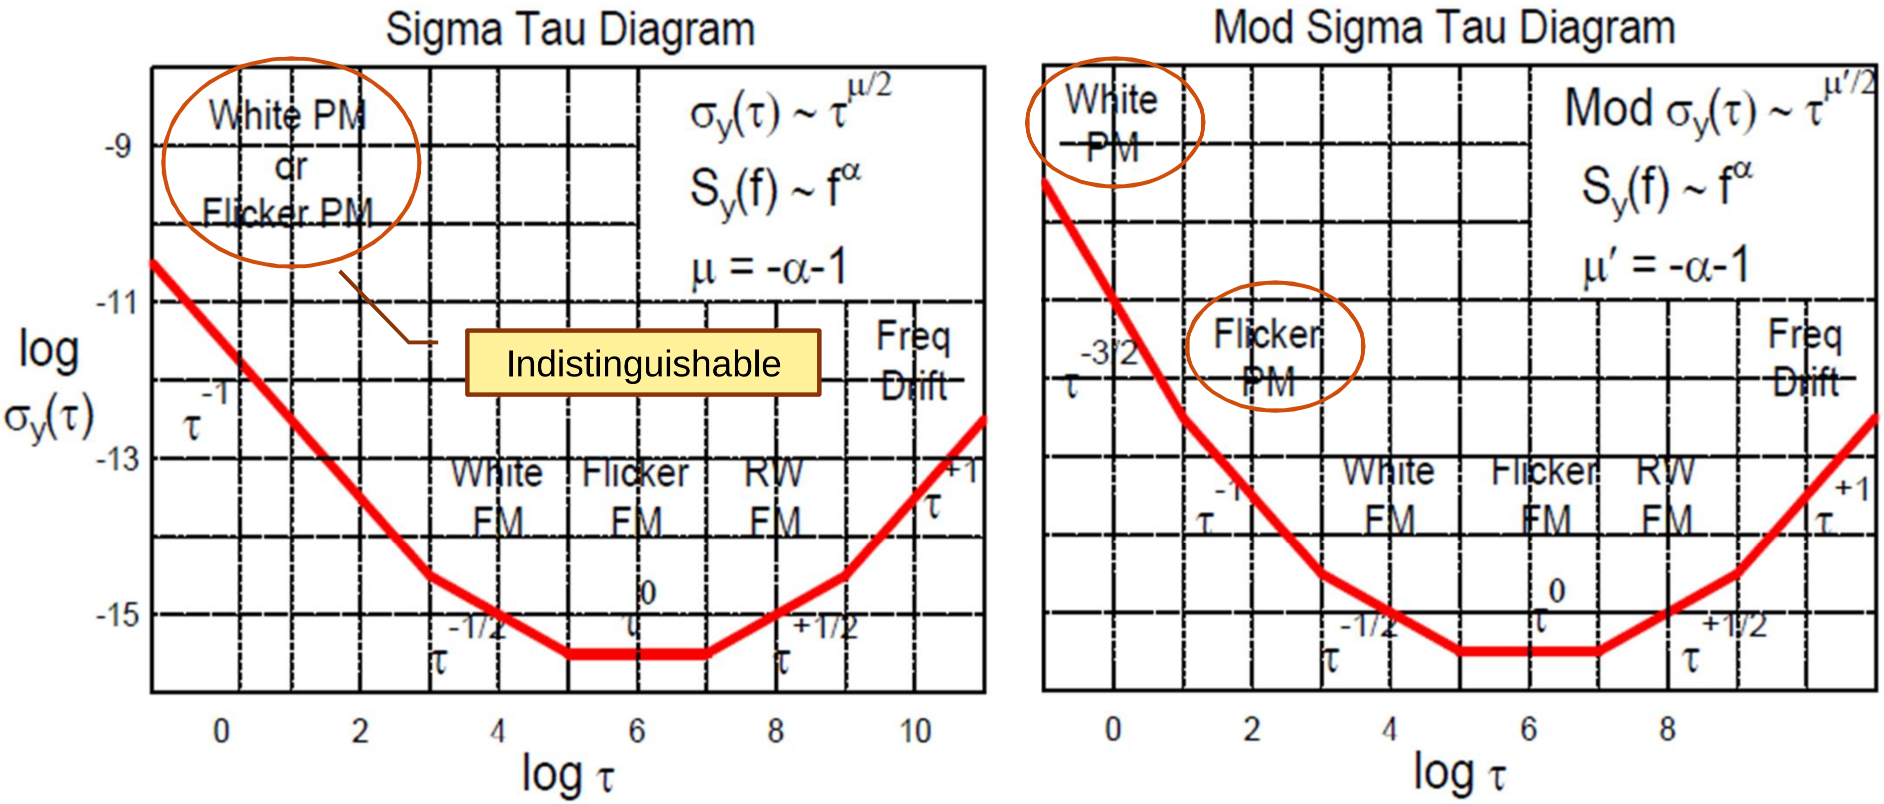

PM = Phase modulaton
FM = Frequency modulation
RW = Random walk

* Std. Deviation (σ) often used instead of Variance (σ^2) in plots

Microsemi

Higher-order Variance Estimation Schemes

- Standard Variance
  - Squared Deviation of sample from mean
  - Donõt useé diverges for many noise types
- Allan Variance
  - Squared Deviation of first back-difference (i.e., discrete-time derivative) of fractional frequency *averages*é
  - Same as second difference of phase *samples*
- Hadamard Variance
  - Same as power-law slopes as Allan Variance but computed from 2nd back-difference of fractional frequency samples (or third difference of phase samples)

- Modified [Allan, Hadamard] Variances
  - Additional phase averaging is done so that White PM and Flicker PM slopes can be distinguished

Averaging Time Interval

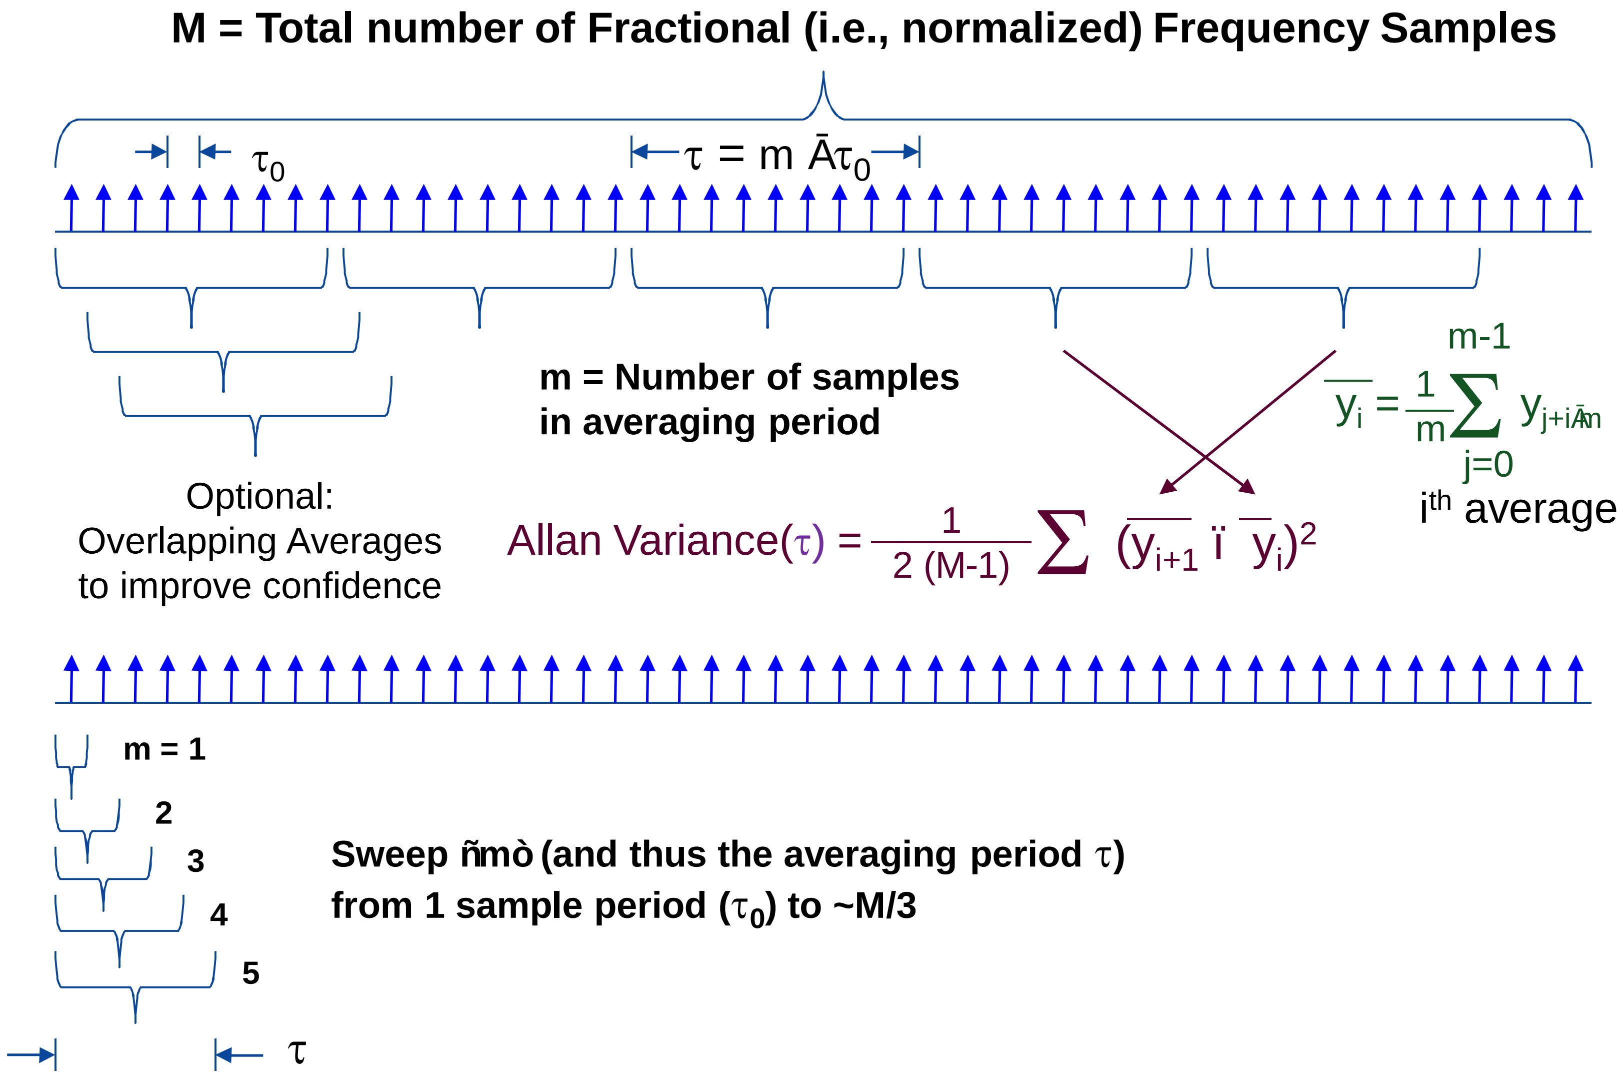

A Few Additional Formulas

N is total number of samples

$$m_{max} := \text{floor}\left(\frac{N}{q}\right)$$

$$\text{for } m = 1 \,..\, m_{max}$$

$$\tau = m\cdot\tau_0$$

Allan_Variance(τ) =

Modified Allan Variance for phase data (q=3)

$$\left[\frac{1}{2\cdot m^2\cdot(m\cdot\tau_0)^2\cdot(N-q\cdot m+1)}\right]\cdot\sum_{j=0}^{(N-q\,m)}\left[\sum_{i=j}^{(j+m-1)}\left(x_{i+2\cdot m}-2\cdot x_{i+m}+x_i\right)\right]^2$$

Overlapping Hadamard Variance for phase data (q=3)

$$\frac{1}{6\cdot(N-q\cdot m)\cdot(\tau_0\cdot m)^2}\cdot\sum_{i=0}^{N-q\cdot m-1}\left(x_{i+3\cdot m}-3\cdot x_{i+2\cdot m}+3\cdot x_{i+m}-x_i\right)^2$$

Overlapping Hadamard Variance for fractional frequency data "Y" (q=3)

$$\frac{1}{6\cdot m^2\cdot(M-q\cdot m+1)\cdot(m\cdot\tau_0)^2}\cdot\sum_{j=0}^{M-q\cdot m}\left[\sum_{i=j}^{j+m+1}\left(y_{i+2\cdot m}-2\cdot y_{i+m}+y_i\right)\right]^2$$

Modified Hadamard Variance for phase data (q=4)

$$\frac{1}{6\cdot m^2\cdot(m\cdot\tau_0)(N-q\cdot m+1)}\cdot\sum_{j=0}^{N-q\cdot m}\left[\sum_{i=j}^{j+m-1}\left(x_i-3\cdot x_{i+m}+3\cdot x_{i+2\cdot m}-x_{i+3\cdot m}\right)\right]^2$$

$$\text{Allan\_Deviation}(\tau) = \sqrt{\text{Allan\_Variance}(\tau)}$$

Test Setup ï SmartFusionÊ Dev. Kit

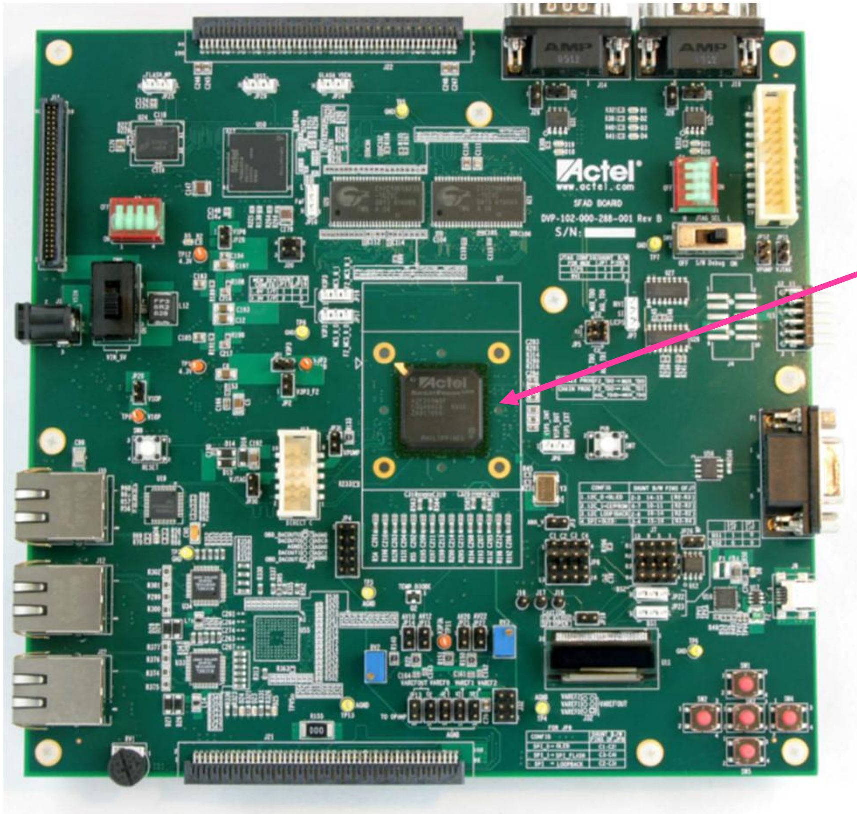

SmartFusion A2F500

FPGA Ring Oscillator (vs. Crystal Osc.)

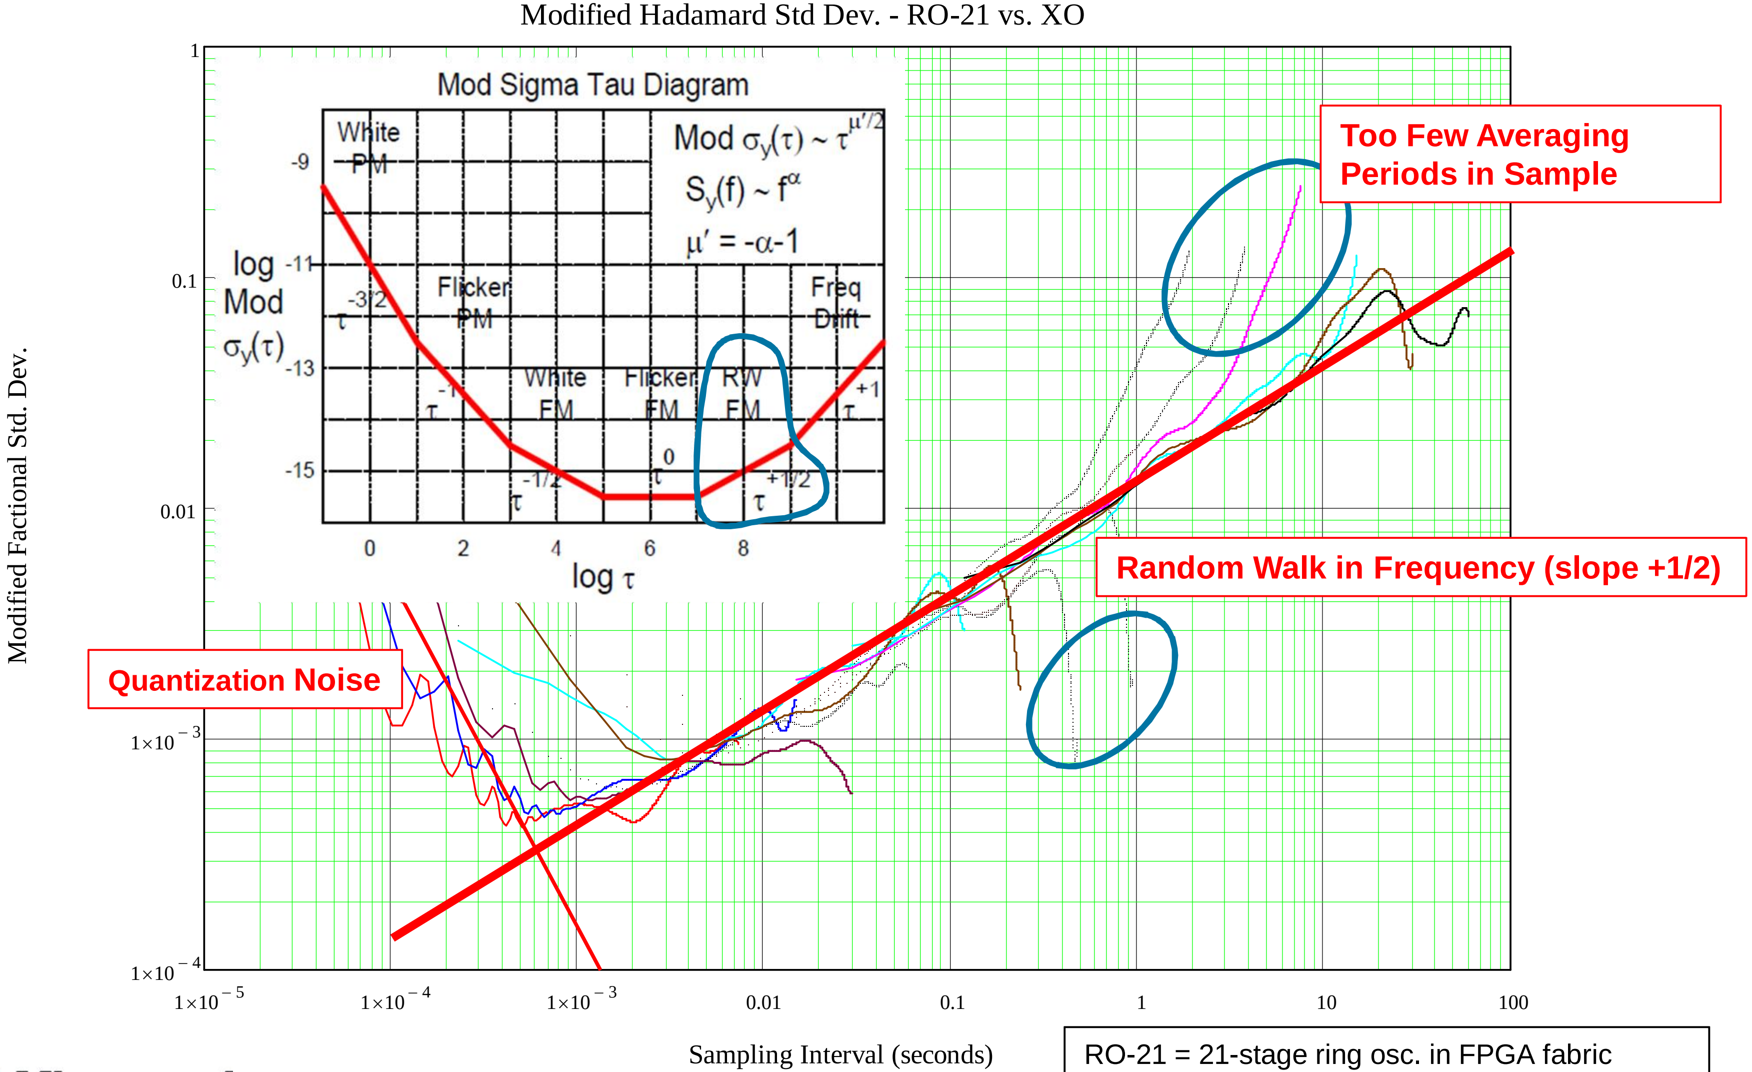

RO-21 = 21-stage ring osc. in FPGA fabric
XO = SmartFusion crystal osc. w/ 50MHz crystal

Microsemi

FPGA R-C Oscillator (vs. Crystal Osc.)

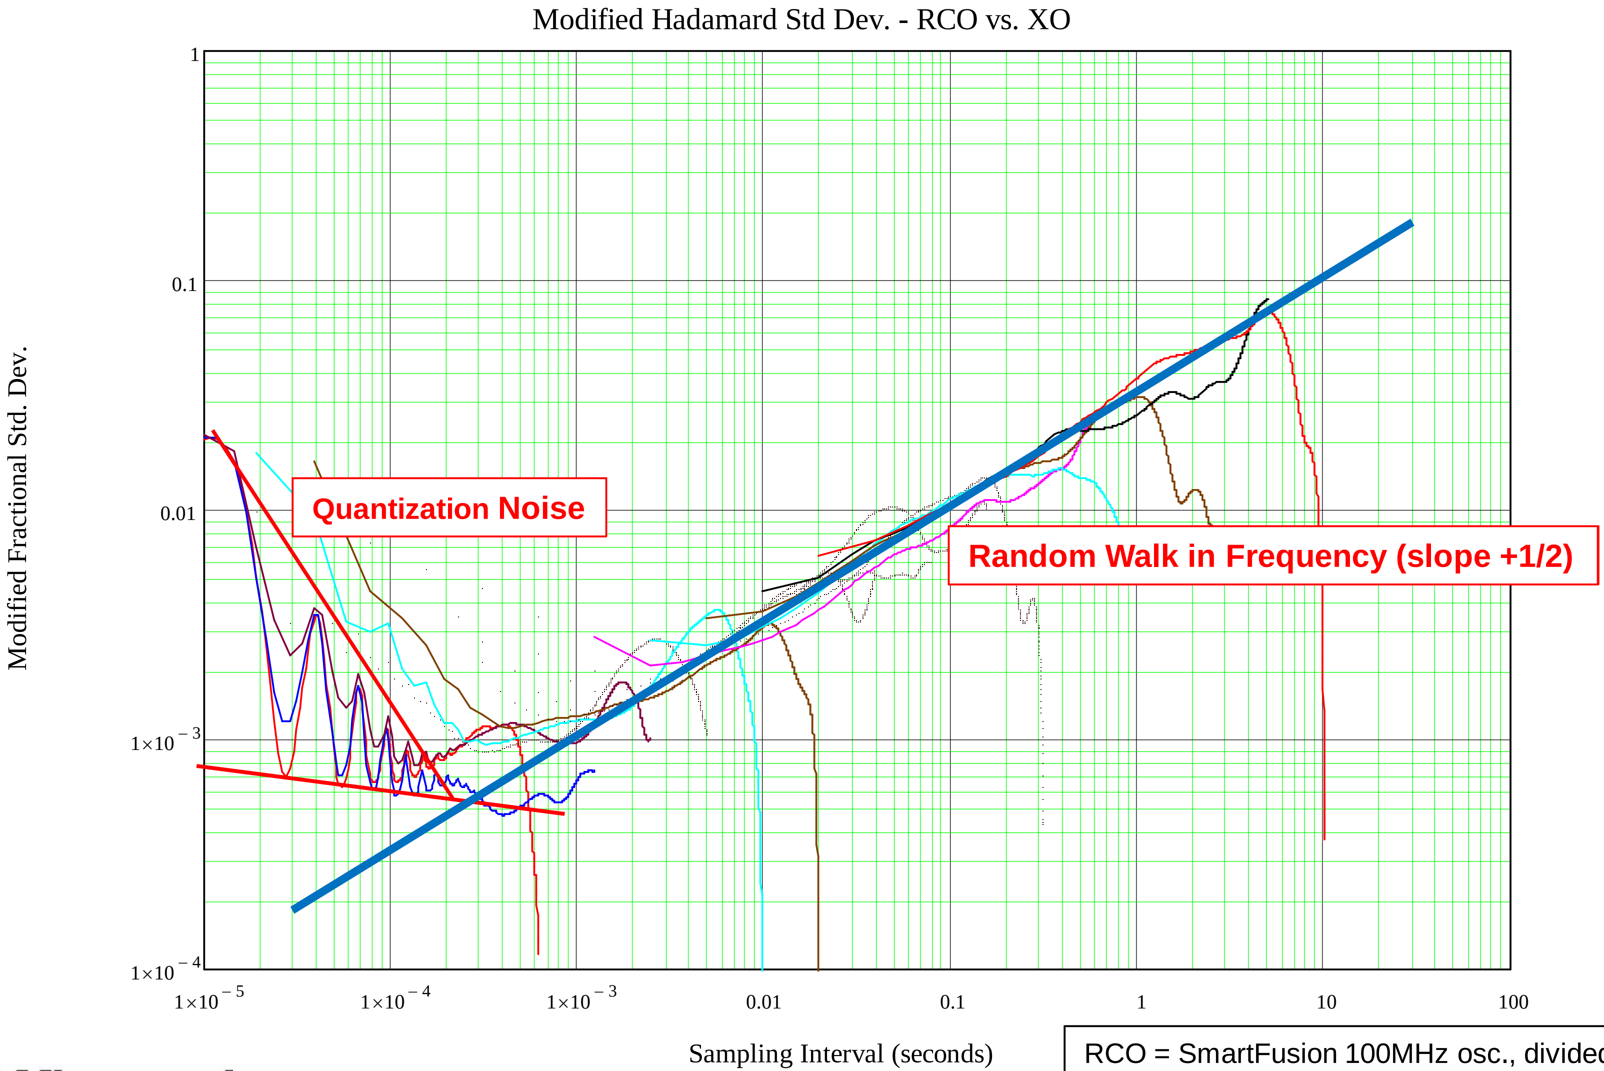

RCO = SmartFusion 100MHz osc., divided by 4
XO ï SmartFusion crystal osc. w/ 50MHz crystal

21-stage Ring Oscillator (vs. 23-stage Ring Osc.)

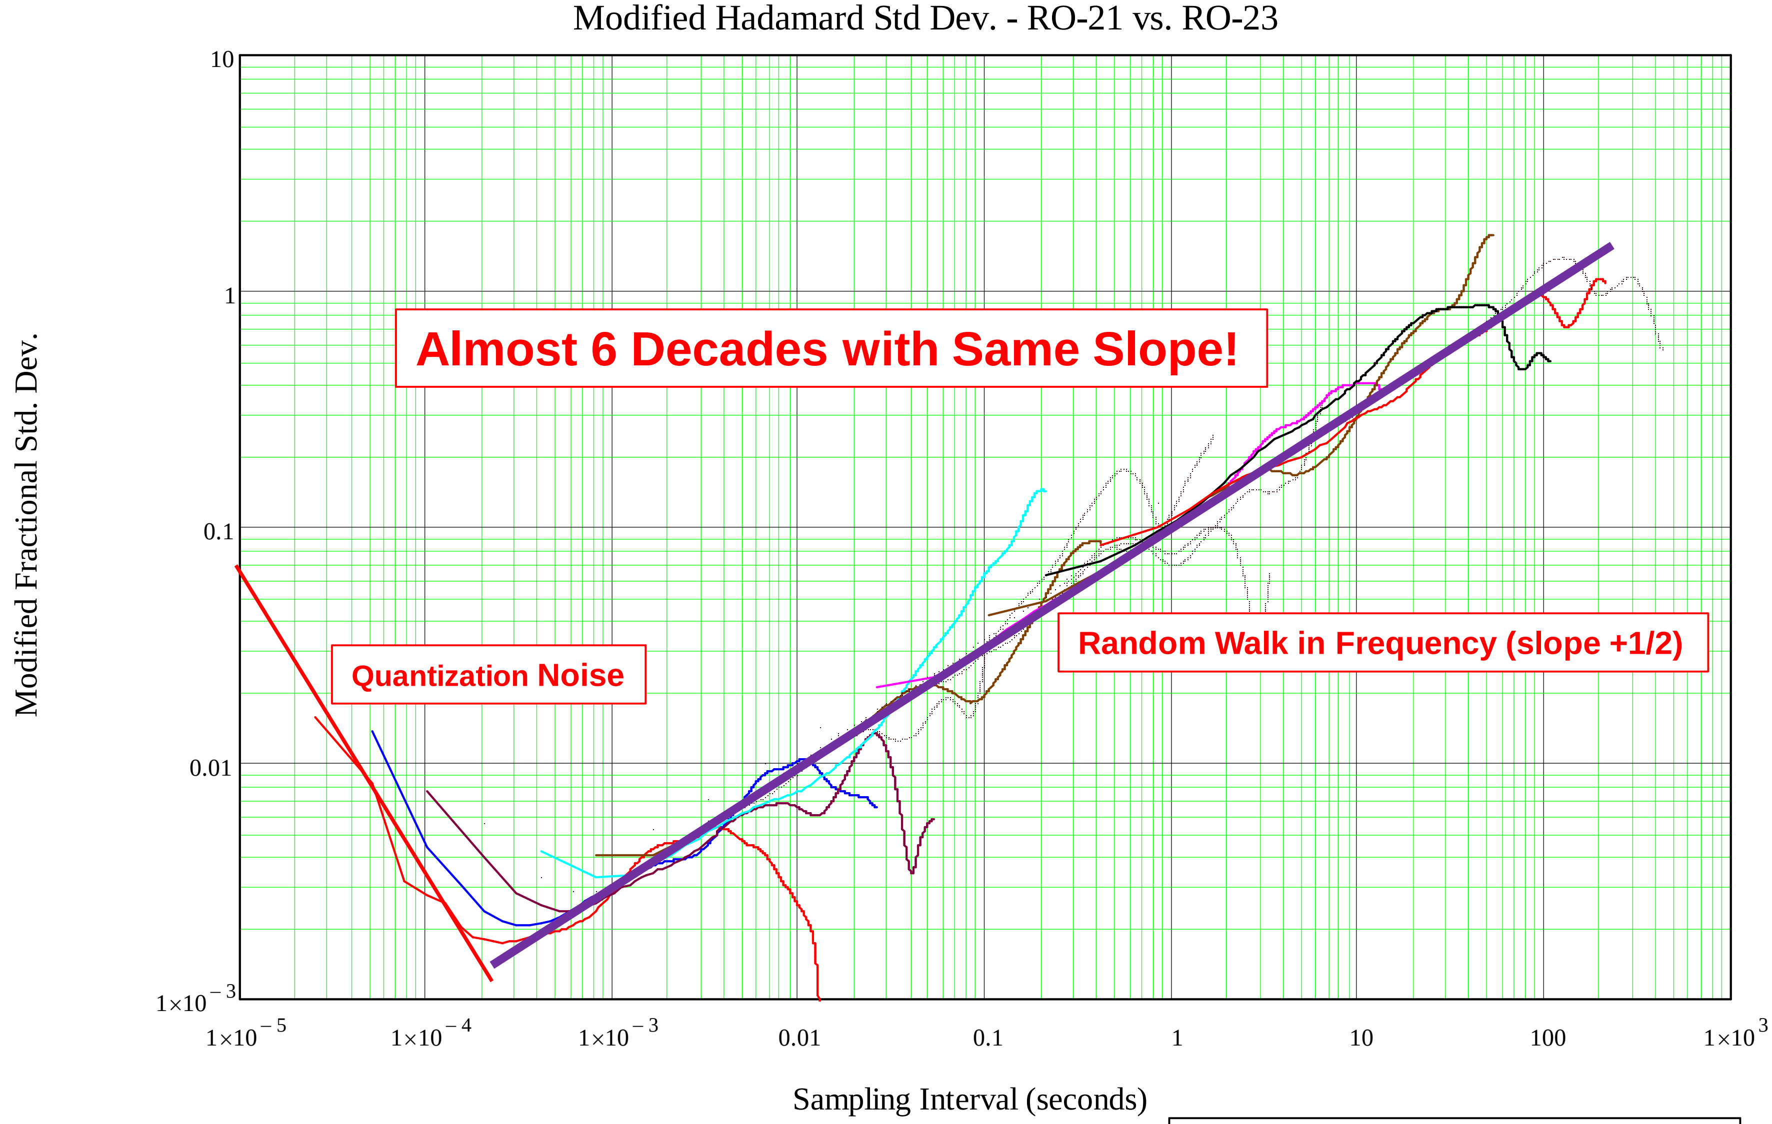

RO-21 = 21-stage ring osc. in FPGA fabric
RO-23 = 23-stage ring osc. in FPGA fabric

Composite Plot

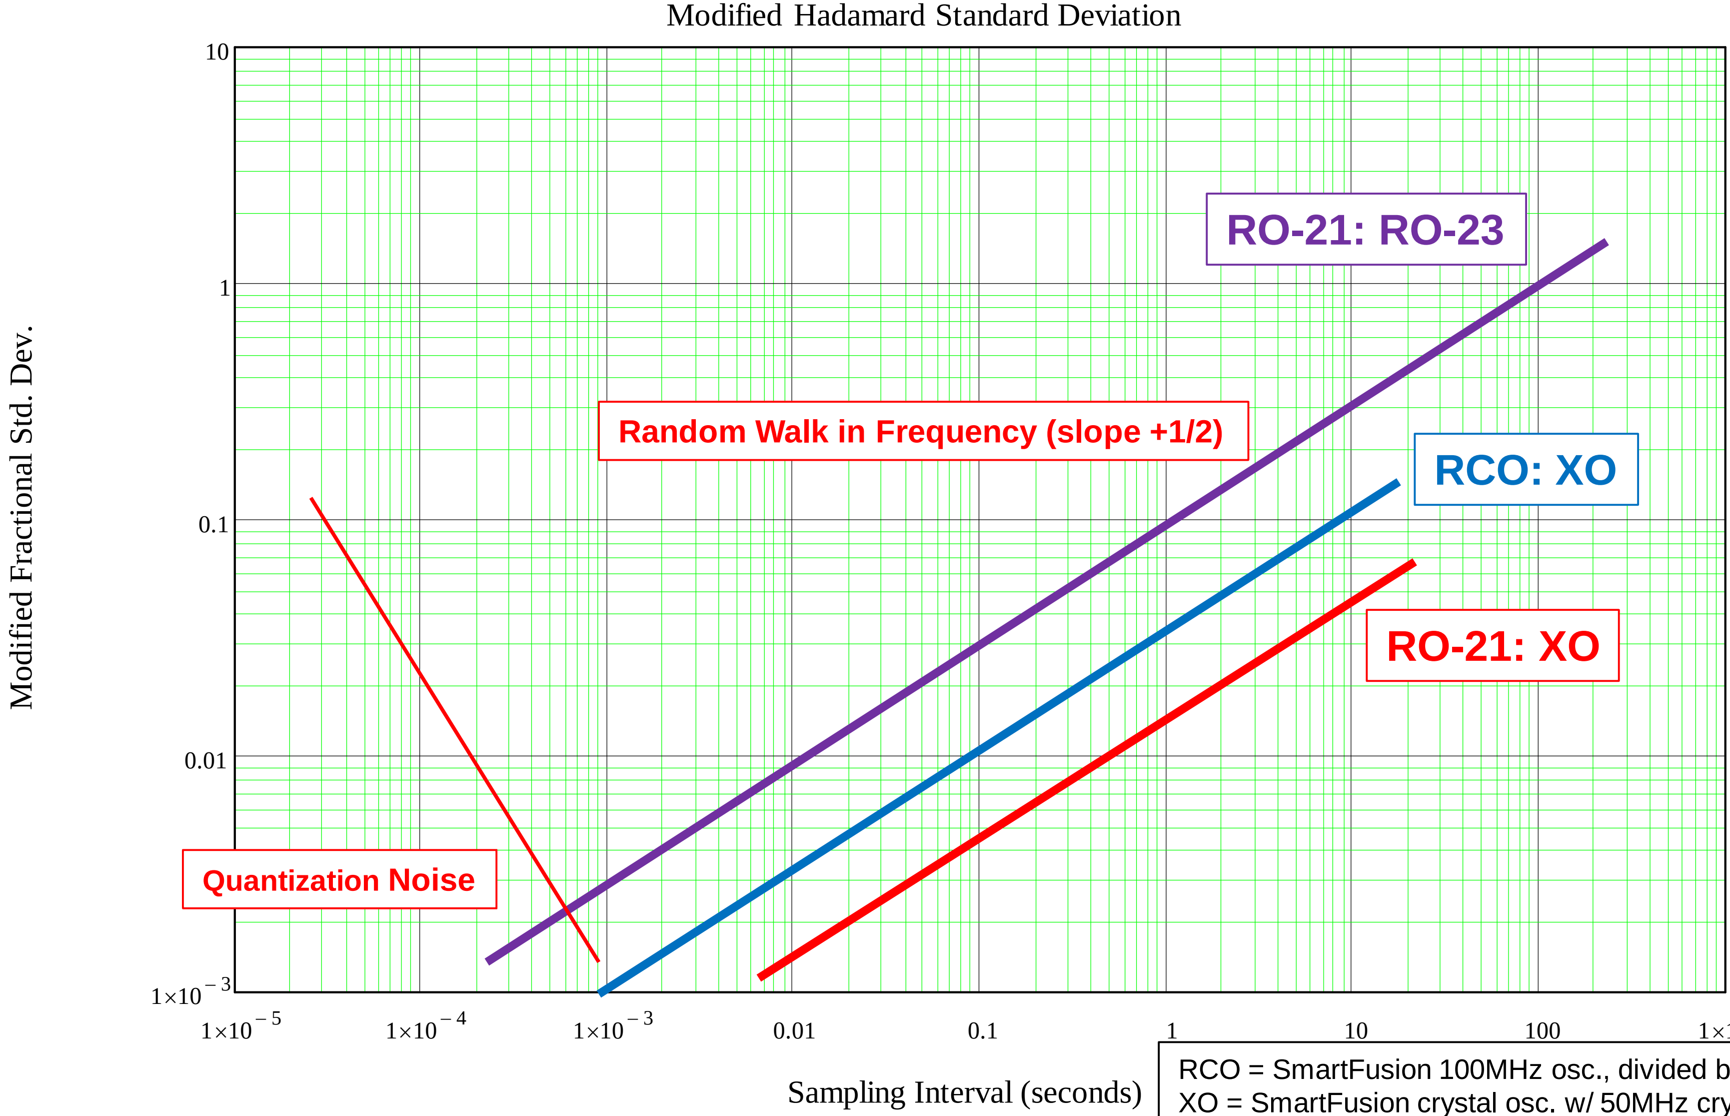

RCO = SmartFusion 100MHz osc., divided by 4
XO = SmartFusion crystal osc. w/ 50MHz crystal
RO-21 = 21-stage ring osc. in FPGA fabric
RO-23 = 23-stage ring osc. in FPGA fabric

RO:XO Time Series (Frequency Samples)

Sampling Frequency: 93.9 Hz

RO: 8.98 MHz
XO: 50MHz
XO-based sampling interval: 532480 counts

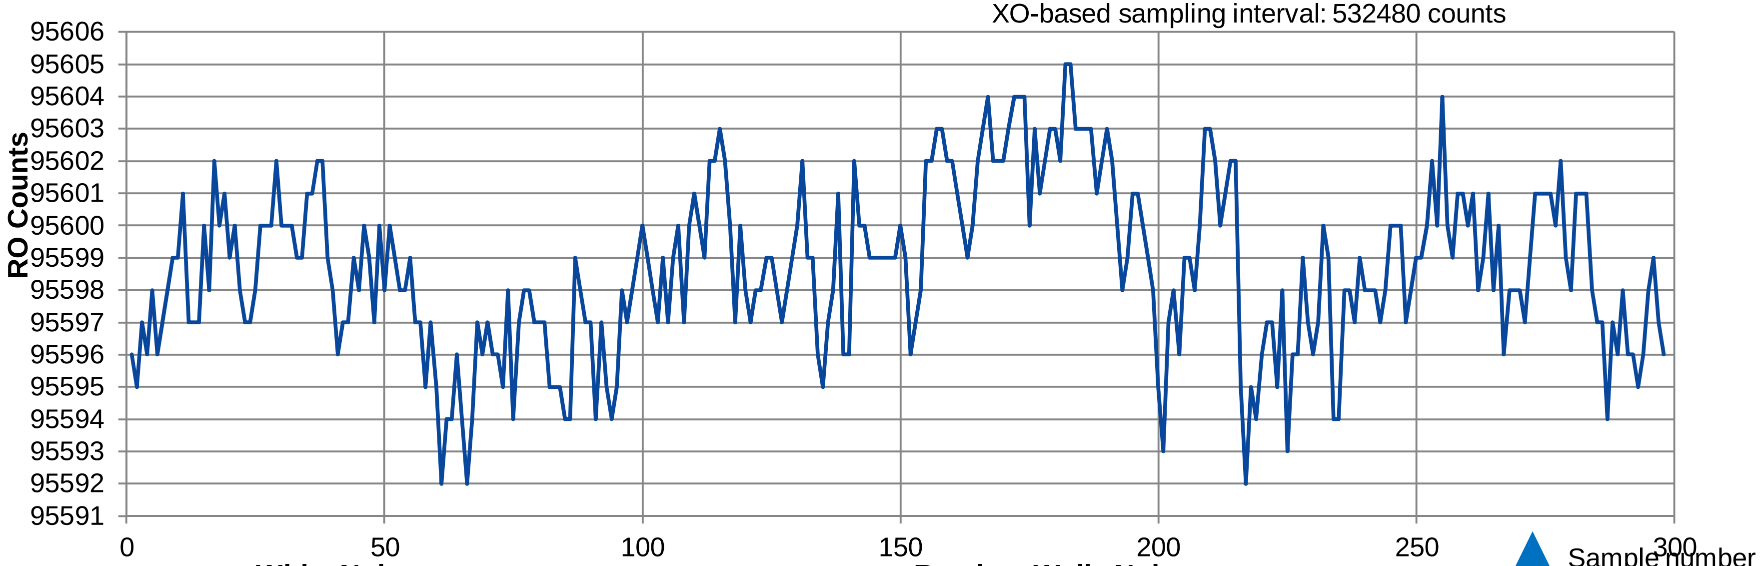

White Noise

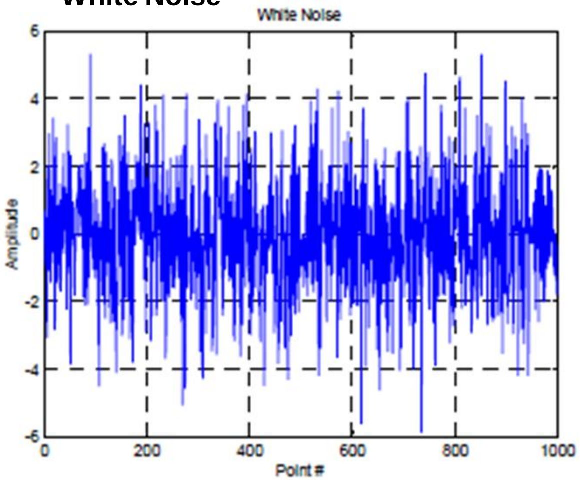

Random Walk Noise

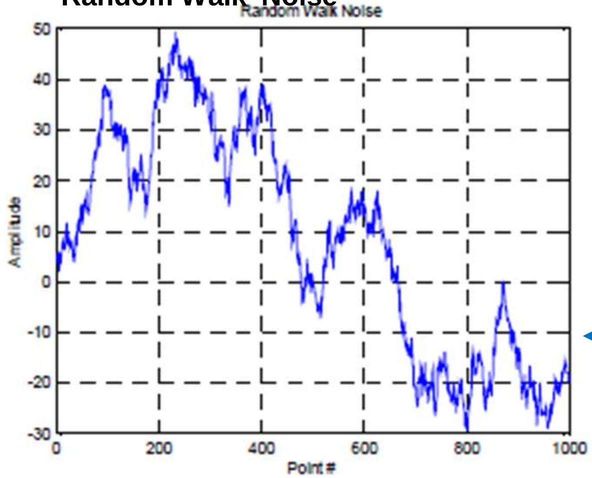

Measured Ring-Osc. Frequency Noise Data

Noise Exemplars

Microsemi

Temperature Sensitivity

- The ring oscillator (RO) showed a strong warm-up characteristic every time it was started. (RCO did not)
  - Å With a 2 to 3 second time-constant
- This was attributed to thermal effects
  - Not likely to cause a problem in a TRNG, but makes analysis more difficult
- To mitigate this effect for the experimental results, the system was started and several buffers full of data discarded, before recording the buffer that was ultimately transferred to the PC, thus allowing the device to thermally stabilize

Frequency Samples
Un-normalized DUT counts

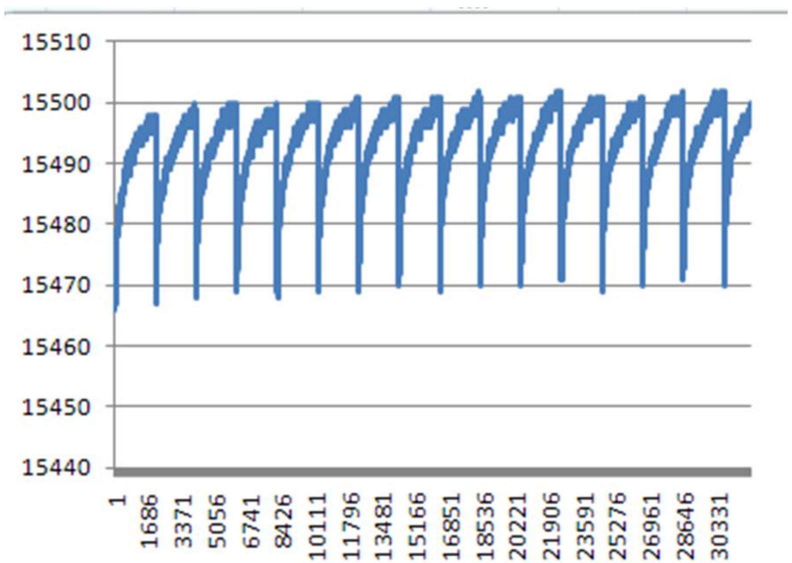

Sample Number

Sample Frequency 542 Hz

Note:
After each 2048-point buffer was filled, the oscillator and frequency counter circuit was put in reset while data was transferred to the PC. The 2048 sample period is clearly visible in the results

Allowing several thermal time-constants before taking data removed this effect

Observations

- All oscillators tested showed a definite ñrandom walk in frequencyò characteristic
  - Note that this is the same as a ñrandom run in phaseò
- The R-C oscillator (RCO) was about twice as noisy as a 21-stage ring oscillator (RO-21) on the same chip, with both referenced to the crystal oscillator (XO)
- When two ring oscillators (RO-21 vs. RO-23) were compared, the noise was roughly six times higher than the RO-21 vs. the XO. This is somewhat surprising
  - If the ROÕs individually experienced common mode noise, it should be reduced when two similar ROÕs are compared
  - The random internal noise of the ROÕs when added, if similar in magnitude, and assuming the noise of the XO is negligible, should increase the RMS noise by a factor of square-root of 2 (not six-fold)
  - We have to conclude that the RO-23 was substantially noiser than the RO-21, but with no explanation forthcoming

Conclusions

- The common perception that Ring Oscillator

  ñflipping times [are] independent and identically distributed,ò

  appears to be wrong, based on these test results
  - ñFlipping timeònoise appears to be dominated by Brownian (1/f^2) circuit noise over a very wide range of frequencies. Brownian noise is random walk noise
  - This is integrated (again) when accumulated into phase, making a random run in phase or, equivalently, a random walk in frequency power-law characteristic
  - Any assumptions about ñindependence,òñmemorylessness,òetc. regarding osc. noise should be carefully validated with experimental data before acceptance

- The Allan Variance time-domain method(s), well known by scientists and engineers in the precision clock and inertial navigation fields, seems less well known amongst cryptologists, whom it may also benefit
  - Low cost to instrument and test
  - Digital signals automatically reject amplitude noise
  - Converges for most power-law noise characteristics, unlike the standard variance which diverges in the presence higher-order power law noise slopes
  - Better suited for low frequency, long time interval stability measurements than frequency-domain methods

Required Reading

NIST Special Publication 1065
Handbook of Frequency Stability Analysis

by W.J Riley

http://tf.nist.gov/general/pdf/2220.pdf