FFT-BASED METHODS FOR SIMULATING FLICKER FM

Charles A. Greenhall[1]
Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Dr., MS 298-100, Pasadena, CA 91109, USA
E-mail: *charles.greenhall@jpl.nasa.gov*

Abstract

Four algorithms for simulating flicker FM phase noise (f^{-3} spectrum) are given, two old and two new. Their Allan deviation and mean square time interval error (MSTIE) are examined. The MSTIE shows that one of the old algorithms has deficient long-term phase deviations on average. The Allan deviation does not reveal this deficiency.

1 INTRODUCTION

By a flicker FM model we mean a stochastic process that has a spectral density (in a sense that can be made precise) that is asymptotically equal to const $/f^3$ as $f \to 0$. Many quartz oscillators exhibit flicker FM phase noise over wide intervals of Fourier frequency, the evidence being Allan deviation plots that are approximately flat over two or more decades of averaging time. Therefore, simulations of systems containing quartz oscillators need to include flicker FM generators. Following are three classes of existing flicker FM generation algorithms. They all have running time of order $N \log N$, where N is the number of points to be generated.

- Filter cascade generators [**1**]. White noise is applied to a set of first-order digital filters in series; the filters are so designed that the output is a stationary process with approximate spectrum f^{-1} over a frequency range whose low end depends on the number of filter stages. One can obtain f^{-3} noise by a cumulative sum operation on the f^{-1} noise. Because these algorithms generate the output sequentially, they take little memory.
- Discrete spectrum (DS) generators. Complex-valued Hermitian white noise (the Fourier transform of real white noise) is generated in the discrete frequency domain, multiplied by $f^{-1/2}$ or $f^{-3/2}$, and transformed back to the time domain. This is a special case of a general method for generating colored noise; the author has no citation for its actual use to generate flicker FM, but Ref. 3 of [**2**] cites a suggestion for its use.

[1]This work was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

- Impulse response (IR) generators, discrete-time analogs of Riemann-Liouville fractional integration (see **[2]** for continuous-time power-law noise models). White noise is convolved with a causal filter that represents fractional summation of order 1/2 or 3/2. The Kasdin-Walter algorithm **[3,4]** discrete convolution in $N \log N$ time by using the FFT.

Our principal aim here is to compare a DS generator, an IR generator, and two new FFT-based flicker FM generators that use the recent method of *circulant embedding* for exact simulation of stationary processes **[5-8]**. The circulant embedding algorithm is given below; one may think of it as the discrete spectrum method with a twist. Two statistical properties are used for the comparison: 1) Allan deviation; 2) a form of mean square time interval error (MSTIE) with an initial phase and frequency removed. Lest this work be merely a discussion of models and algorithms, the comparison also includes the phase residuals of two precision quartz oscillators that were chosen for flatness of their Allan deviations. The MSTIE (but not Allan deviation) exposes a deficiency of the IR generator, the same deficiency that a filter cascade generator has if it is not properly initialized **[9]**: the generator's long-term phase excursions are too small on average. Nevertheless, as Schmidt also found **[10]**, one can avoid this deficiency by generating twice as many points as needed and using only the second half.

We will be working with discrete-time phase models x_n, $-\infty < n < \infty$, with sample period 1, normalized so that their *two-sided* spectral densities $S_x(f)$, $|f| \leq 1/2$, are asymptotic to $|2\pi f|^{-3}$ as $f \to 0$. To simulate samples $x(n\tau_0)$ of flicker FM phase (time) $x(t)$ with one-sided frequency spectrum $S_y^+(f) = h_{-1} f^{-1}$ and Allan deviation $\sqrt{h_{-1} \ln 4}$, multiply x_n by $\sqrt{\pi h_{-1}} \tau_0$.

2 TWO FLICKER FM PHASE MODELS

Before defining the four generators, we define two flicker FM models, against which the generators can be compared. Although each model for phase x_n is a nonstationary stochastic process, its second increment $\Delta^2 x_n = x_n - 2x_{n-1} + x_{n-2}$ is a stationary, mean-zero, Gaussian process. The definition of x_n is ambiguous in that any constant phase and frequency can be added to it; to pin down x_n exactly, we can fix two values x_a and x_b. The two models differ mainly in their spectral densities near the Nyquist frequency $f = 1/2$ (see Figure 1).

2.1 FD(3/2) (Fractional Difference) Model

For each value of the real parameter δ, there is a process called FD(δ) **[11,12]** with spectral density $|2 \sin \pi f|^{-2\delta}$ (in a sense to be explained for $\delta = 3/2$). This family has the convenient property that if x_n is an FD(δ), then Δx_n is an FD($\delta - 1$). (The frequency response of the difference operator Δ is $|2 \sin \pi f|$.) In particular, FD($-1/2$) is defined as a stationary, mean-zero, Gaussian process z_n with spectral density $|2 \sin \pi f|$. For its autocovariance (ACV) sequence, $s_{z,n} = \mathrm{E} z_j z_{j+n}$, we have

$$s_{z,n} = \int_0^1 e^{i2\pi fn} (2 \sin \pi f)\, df = \frac{1}{\pi \left(\frac{1}{4} - n^2\right)}. \tag{1}$$

By definition, x_n *is an FD*$(3/2)$ *process if* $\Delta^2 x_n$ *is an FD*$(-1/2)$ *process*. Then x_n is a process with stationary second increments, and $S_x(f) = |2\sin\pi f|^{-3}$ in the following sense: if H is a finite moving-average filter that contains Δ^2 as a factor, then Hx_n is stationary and $S_{Hx}(f) = \left|H\left(e^{-i2\pi f}\right)\right|^2 S_x(f)$, where $H(z)$ is the z-transform of the impulse response.

It can be proved that an FD(3/2) process x_n can be represented in the following way:

$$\begin{aligned} &x_n - \left(1+\frac{n}{n_1}\right)x_0 + \frac{n}{n_1}x_{-n_1} \\ &= \sum_{j=1}^{n} h_{n-j}u_j + \sum_{j=-\infty}^{0}\left[h_{n-j} - \left(1+\frac{n}{n_1}\right)h_{-j} + \frac{n}{n_1}h_{-n_1-j}\right]u_j, \quad n \geq 1, \end{aligned} \tag{2}$$

where n_1 is a positive integer, u_n is a standard white noise sequence (independent Gaussians with mean 0 and variance 1), and h_n is defined by the power series $(1-z)^{-3/2} = \sum_{n=0}^{\infty} h_n z^n$; thus, $h_n = (-1)^n \binom{-3/2}{n}$, and $h_n = 0$ for $n < 0$ by convention. One can show that $h_n \sim 2\sqrt{n/\pi}$ as $n \to \infty$.

2.2 Sampled PPL (Pure Power Law) Model

Starting with a continuous-time process $x(t)$ with stationary second increments and spectral density $|2\pi f|^{-3}$ for all nonzero real f, we sample it at the integers to get a discrete-time process $x(n)$. Its second increment, $z(n) = \Delta^2 x(n)$, is stationary and has ACV

$$s_z(n) = s_x(n+2) - 4s_x(n+1) + 6s_x(n) - 4s_x(n-1) + s_x(n-2), \tag{3}$$

where $s_x(t)$ is the generalized ACV **[13,14]** of $x(t)$:

$$s_x(t) = \frac{1}{2\pi}t^2 \ln|t| \quad \text{for } t \neq 0, \quad s_x(0) = 0. \tag{4}$$

The spectral density of the sampled process is not $|2\pi f|^{-3}$, but

$$\sum_{k=-\infty}^{\infty} |2\pi(f+k)|^{-3}, \quad |f| \leq 1/2 \tag{5}$$

(see Figure 2).

3 TWO GENERAL ALGORITHMS

Three of the generators under discussion can be quickly given in terms of two algorithms of wider utility; the second algorithm uses the first. With small changes, these descriptions follow Percival and Walden **[5]**. We use a standard complex discrete Fourier transform (called the FFT from now on) and its inverse, defined for M-point sequences by

$$A_k = \sum_{n=0}^{M-1} a_n \exp\left(-i2\pi kn/M\right), \quad a_n = \frac{1}{M}\sum_{k=0}^{M-1} A_k \exp\left(i2\pi kn/M\right),$$

respectively.

3.1 Discrete Spectrum Algorithm

Purpose: Generate values of a real stationary Gaussian process with a desired discrete spectrum.

Inputs: Nonnegative numbers $S_0, \ldots, S_N$, where N is a power of 2 and S_k is the desired two-sided spectral density at frequency $f_k = k/(2N)$.

Outputs: Gaussian random variables $z_0, \ldots, z_N$ such that $\mathrm{E}z_n = 0$,

$$\mathrm{E}z_m z_n = \frac{1}{2N}\sum_{k=1-N}^{N} S_{|k|} \exp\left(i2\pi f_k (n-m)\right).$$

Procedure:

Generate $U_0, U_1, \ldots, U_N, V_1, \ldots, V_{N-1}$ as independent standard Gaussians.

Let $Z_0 = \sqrt{S_0}U_0$, $Z_N = \sqrt{S_N}U_N$.

Let $Z_k = \sqrt{S_k/2}\,(U_k + iV_k)$, $Z_{2N-k} = Z_k^*$ for $k = 1$ to $N-1$.

Let the real-valued sequence $z_0, \ldots, z_{2N-1}$ be $\sqrt{2N}$ times the inverse $2N$-point FFT of the Hermitian sequence $Z_0, \ldots, Z_{2N-1}$.

Keep the values $z_0, \ldots, z_N$ (or any $N+1$ consecutive values).

3.2 Circulant Embedding Algorithm

Purpose: Generate values of a real stationary Gaussian process with a given autocovariance (ACV).

Inputs: Real numbers $s_0, \ldots, s_N$, the desired ACV up to lag N, where N is a power of 2.

Outputs: Gaussian random variables $z_0, \ldots, z_N$ such that $\mathrm{E}z_n = 0$, $\mathrm{E}z_m z_n = s_{n-m}$ for $0 \le m \le n \le N$.

Procedure:

Let $\tilde{s}_n = s_n$ for $n = 0$ to N, $\tilde{s}_{2N-n} = s_n$ for $n = 1$ to $N-1$ (extension of s_n by reflection).

Remark: "Circulant" refers to the covariance matrix that corresponds to $\tilde{s}_n$.

Let the real-valued sequence $\tilde{S}_0, \ldots, \tilde{S}_{2N-1}$ be the $2N$-point FFT of $\tilde{s}_0, \ldots, \tilde{s}_{2N-1}$.

If any $\tilde{S}_k < 0$, the method *fails.* (This means that the extended circular sequence is not positive-definite.)

Use $\tilde{S}_0, \ldots, \tilde{S}_N$ in the discrete spectrum algorithm to generate the z_n.

4 FOUR FLICKER FM GENERATORS

All these generators produce approximately N phase values x_n, where N is a power of 2. The first two are approximate; they do not simulate either target model exactly. The last two give exact simulations of the two target models.

4.1 DS – Discrete Spectrum

Let $f_k = k/(2N)$. Run the discrete spectrum algorithm with input $S_0 = 0$, $S_k = (2\pi f_k)^{-3}$, $k = 1, \ldots, N$, and output $x_0, \ldots, x_N$. (One could also set $S_k = (2 \sin \pi f_k)^{-3}$ to approximate the FD(3/2) model more closely.)

4.2 IR – Impulse Response

This generator is an approximate simulation method for the FD(3/2) model. Its output is given by the formula

$$x_n = \sum_{j=1}^{n} h_{n-j} u_j, \quad 1 \le n \le N, \tag{6}$$

where h_n is defined after (2), and $u_1, \ldots, u_N$ are independent standard Gaussians. The convolution is carried out by zero-padding the sequences to length $2N$, Fourier transforming them, multiplying the transformed sequences, and inverse transforming the result. For details, see **[3]**. Observe that (6) is just one part of (2).

4.3 FD – Fractional Difference

This generator is an exact simulation of $N+3$ values x_n of the FD(3/2) model. Run the circulant embedding algorithm using the ACV (1) for the input $s_0, \ldots, s_N$, and $z_0, \ldots, z_N$ as output. It can be proved that the algorithm succeeds. (The ACV satisfies Craigmile's criterion [**8**].) This produces an exact realization of $N+1$ values of FD(−1/2). Then perform two cumulative summations:

$$y_n = y_{n-1} + z_{n-1} \quad \text{for } n = 1 \text{ to } N+1,$$
$$x_n = x_{n-1} + y_{n-1} \quad \text{for } n = 1 \text{ to } N+2.$$

The initial values y_0 and x_0 are arbitrary, and may be set to zero.

4.4 PPL – Pure Power Law

This generator is an exact simulation of $N + 3$ values of the sampled PPL model. It is identical to the FD generator just described except that the ACV (3) is used in place of (1). Again, it can be proved that the algorithm succeeds. There is one complication: to avoid catastrophic roundoff error in (3), use the asymptotic approximation

$$s_z(n) \sim -\frac{1}{\pi n^2}\left(1 + \frac{1}{n^2} + \frac{3}{2n^4}\right) \tag{7}$$

in place of (3) whenever $n \geq 35$.

5 COMPARISONS

We compare the four flicker FM generators with each other and with the phase residuals of two quartz oscillators (Oscilloquartz and C-MAC) that were compared once per second against hydrogen masers. The test runs were chosen for flatness of Allan deviation between 1 and 1,000 seconds[2].

Figure 1 shows a sample output of the four generators with $N = 1024$. Also shown are the first 1,025 phase residuals of the quartz oscillators, scaled up as explained in the next section. The exact generators are both initialized so that $x_0 = x_1 = 0$. The IR output starts with a small nonzero value of x_1. The DS output, which realizes a stationary process, has no special initial value.

Figure 2 shows the spectral density of the two target models along with the discrete spectrum of the DS generator for $N = 32$. The spectral densities, multiplied by $(2\pi f)^3$, are plotted on a linear scale against f. As the next section shows, the rise in the sampled PPL spectrum (5) near the Nyquist frequency is just right to make its Allan deviation exactly flat for all integral τ. The DS spectrum actually has too little high-frequency power for this purpose, the FD spectrum too much. For many purposes these high-frequency deviations are of little concern.

5.1 Allan Deviation

Figure 3 shows the theoretical Allan deviation (lines) and the square root of measured Allan variance (small dots), averaged over 10,000 trials, for the four generators with $N = 1,024$. Also shown are the Allan deviations of the quartz oscillators (symbols) as estimated from their test runs, which lasted about 45,000 s for the Oscilloquartz, 129,000 s for the C-MAC. The oscillator results were normalized so that $\sigma_y(64\,\text{s}) = \sqrt{(\ln 4)/\pi} \doteq 0.664$, the theoretical value assumed by the PPL model for all τ. The σ axis has an expanded linear scale to bring out the differences among the plots. The IR generation was actually performed with $N = 2,048$; IR1 refers to the first half of the generated sequence, which is equivalent to IR with $N = 1,024$; IR2 refers to the second half.

[2]Thanks to Albert Kirk for making these tests available.

The Allan deviation appears to confirm that all generators behave like flicker FM. The PPL line is exactly flat. For the other generators, we see the expected minor deviations from flatness for small τ and an insignificant droop by DS and IR1 for large τ.

5.2 Two-Point MSTIE

Suppose that the phase $x(t)$ of a clock is measured at times $t_0 - \tau_1$ and t_0. For the purpose of this discussion, the mean square time interval error (MSTIE) of a clock after a delay τ is defined by

$$\text{MSTIE}(\tau, \tau_1) = \text{E}\left[x(t_0+\tau) - \left(1+\frac{\tau}{\tau_1}\right)x(t_0) + \frac{\tau}{\tau_1}x(t_0-\tau_1)\right]^2, \tag{8}$$

which is the mean square error of linear extrapolation from the two phase measurements. This is independent of t_0 for all processes $x(t)$ with stationary second increments, including all models and generators discussed here except IR. Although the method of phase calibration is crude, this measure serves the purpose of showing how the variance of the long-term phase deviations grows as we go farther and farther from a fixed calibration interval. (See **[15]** for a discussion of more sophisticated calibrations.)

Figure 4 plots theoretical and measured values (averaged over 10,000 trials) of $\text{MSTIE}(\tau, 10)/\tau^2$ against τ for the flicker FM generators. For IR1, $t_0 = 1$; for IR2, $t_0 = 1,025$. The MSTIE for the quartz oscillators (normalized as before) was measured by averaging the squared extrapolation error over all available t_0, just as one does when estimating Allan variance. All the curves except IR1 show the same asymptotic $\pi^{-1}[1+\ln(\tau/\tau_1)]$ behavior that has been calculated for the PPL model [**?**], with a tiny droop at the largest τ for the DS generator. The IR2 curve cannot be distinguished from the FD curve, but the IR1 curve falls off significantly from the others as τ increases.

6 CONCLUSIONS

According to the tests that we have applied, the DS, FD, and PPL generators are good fits to the phase noise of a quartz oscillator in its flicker FM range, whereas the IR generator has significantly smaller long-term mean square phase deviations. Although the Allan deviations of the Oscilloquartz and C-MAC oscillators (Figure 3) rise as τ increases beyond 64 s, this rise is not enough to explain why the oscillator MSTIE points (Figure 4) tend to line up with the curves for DS, FD, PPL, and IR2 (the "second-half" modification of IR), but not with the IR1 curve.

The deficiency of the IR generator is exposed by the two-point MSTIE, but not by the Allan deviation. Flatness of Allan deviation is a necessary criterion for a flicker FM generator, but it is by no means sufficient.

The FD and PPL generators are exact simulators of their nonstationary target models. Even though their outputs live on a finite time interval, they contain contributions from arbitrarily low Fourier frequencies. The DS generator output, though it is only a stationary process, still behaves substantially like a nonstationary f^{-3} process on its design interval. We may regard the IR generator as an inaccurate simulator of the FD(3/2) model. Fortunately, one can still get an accurate FD(3/2) simulation from IR by generating twice as many points as one needs and keeping only the second half. In this case, one has to use FFTs of size $4N$ to generate N points. (One can also improve the accuracy of the DS generator by using an FFT size of $4N$ instead of $2N$.)

The IR generator is deficient because it neglects the past of the FD(3/2) process; the same is true for the fractional integration models in **[2]** and the filter cascade generator with zero initial state **[9]**. The IR output (6) is just the first term of the right side of the FD(3/2) formula (2), whose other terms represent the effect of the random shocks u_j from the past, $j \leq 0$, on the future value x_n. It is one thing to tie the present to zero, as we often do; it is another thing to neglect the past of these long-memory processes.

Finally, despite being able to produce models and generators that have the right sort of behavior, the author does not pretend to understand the physical mechanisms behind flicker noise.

7 REFERENCES

[1] J. A. Barnes and S. Jarvis, Jr., 1971, *"Efficient numerical and analog modeling of flicker noise processes,"* **NBS Technical Note, 604** (National Institute of Standards and Technology, Boulder, Colorado).

[2] J. A. Barnes, and D. W. Allan, 1966, *"A statistical model of flicker noise,"* **Proceedings of the IEEE, 54**, 176-178.

[3] N. J. Kasdin and T. Walter, *"Discrete simulation of power law noise,"* Proceedings of the 1992 IEEE Frequency Control Symposium, 27-29 May 1992, Hershey, Pennsylvania, USA (IEEE Publication 92CH3083-3), pp. 274-283.

[4] N. J. Kasdin, 1995, *"Discrete simulation of colored noise and stochastic processes and $1/f^{\alpha}$ power law noise generation,"* **Proceedings of the IEEE, 83**, 802-827.

[5] D. B. Percival and A. T. Walden, 2000, **Wavelet Methods for Time Series Analysis**, Sec. 7.8 (Cambridge University Press, UK).

[6] C. R. Dietrich and G. N. Newsam, 1997, *"Fast and exact simulation of stationary Gaussian processes through circulant embedding of the covariance matrix,"* **SIAM Journal of Scientific Computing, 18**, 1088-1107.

[7] T. Gneiting, 2000, *"Power-law correlations, related models for long-range dependence and their simulation,"* **Journal of Applied Probability, 37**, 1104-1109.

[8] P. F. Craigmile, 2002, *"Simulating a class of stationary Gaussian processes using the Davies–Harte algorithm, with application to long memory processes,"* **Journal of Time Series Analysis**, accepted for publication.

[9] C. A. Greenhall, 1986, *"Initializing a flicker-noise generator,"* **IEEE Transactions on Instrumentation and Measurement, IM-35**, 222-224.

[10] L. S. Schmidt, 2001, *Estimation in the Presence of Fractionally Integrated Noise: an Application to Atomic Timescales,* Ph.D. dissertation, American University, Washington, D.C.

[11] C. W. J. Granger and R. Joyeux, 1980, *"An introduction to long-memory time series models and fractional differencing,"* **Journal of Time Series Analysis, 1**, 15-29.

[12] J. R. M. Hosking, 1981, *"Fractional differencing,"* **Biometrika, 68**, 165-176.

[13] C. A. Greenhall, 1983, *"A structure function representation theorem with applications to frequency stability estimation,"* **IEEE Transactions on Instrumentation and Measurement, IM-32**, 364-370.

[14] C. A. Greenhall, 1999, *"The generalized autocovariance: a tool for clock noise statistics,"* TMO Progress Report 42-137, Jet Propulsion Laboratory, May 1999; online at *http://tmo.jpl.nasa.gov/tmo/progress_report/42-137/137J.pdf*

[15] J. Delporte, F. Vernotte, M. Brunet, and T. Tournier, 2001, *"Modelization and extrapolation of time deviation of USO and atomic clocks in GNSS-2 context,"* in Proceedings of the 32nd Annnual Precise Time and Time Interval (PTTI) Systems and Applications Meeting, 28-30 November 2000, Reston, Virginia, USA (U.S. Naval Observatory, Washington, D.C.), pp. 45-55.

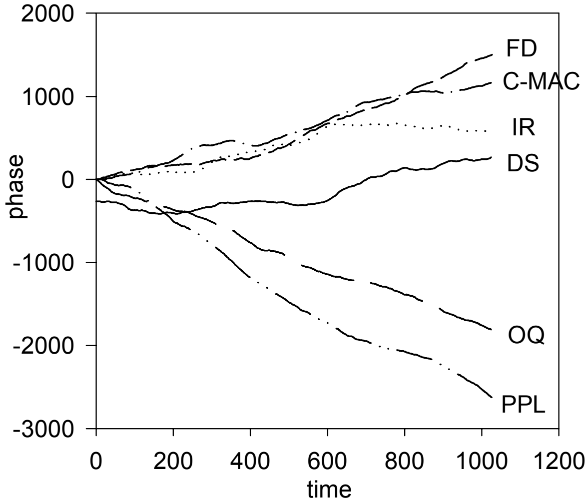

Fig. 1. Phase samples from flicker FM generators and quartz oscillators. Generators: DS = discrete spectrum, IR = impulse response, FD = fractional difference, PPL = pure power law. Oscillators (normalized): OQ = Oscilloquartz, C-MAC.

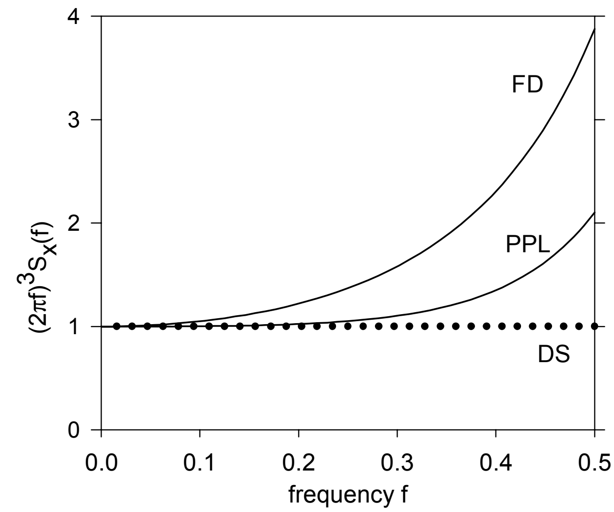

Fig. 2. Spectra of the DS generator (N = 32), the FD(3/2) model, and the sampled PPL model. The spectral densities are multiplied by $(2\pi f)^3$.

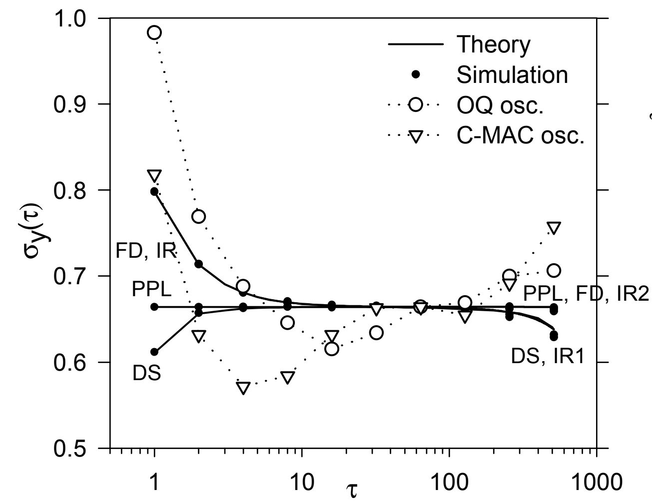

Fig. 3. Allan deviation of flicker FM generators and quartz oscillators (normalized). IR1 = first half of IR output with N = 2048, IR2 = second half.

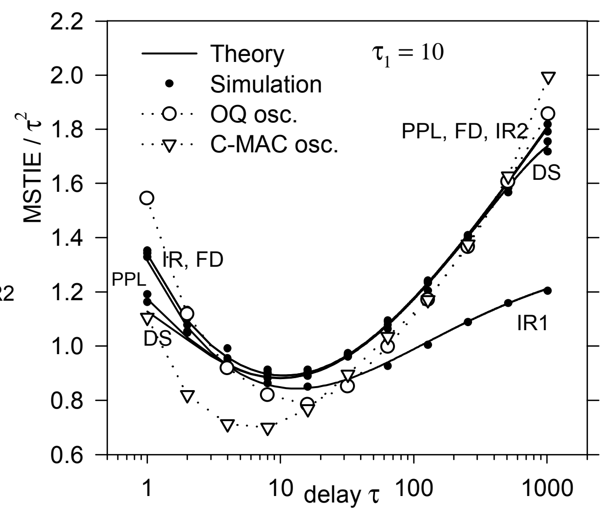

Fig. 4. Mean square time interval error, divided by τ^2, for delay τ and calibration period $\tau_1 = 10$. All the flicker FM generators except IR1 behave like the normalized oscillators for large τ.

QUESTIONS AND ANSWERS

DAVE HOWE (National Institute of Standards and Technology): Chuck, what is the rationale for reflecting the autocorrelation function and then extending it, the embedded circulant approach? Because, it looks like what you are assuming is that the data that you have is a piece of something that is periodic.

CHARLES GREENHALL: You are only trying to simulate N plus 1 points exactly. What it says is that there is a two-end periodic stationary process whose joint distribution is exactly equal to N plus 1 successive values of the stationary process you are trying to simulate.

HOWE: My question is how do you know a priori that this is a noisy quasi-periodic phenomenon? Why would you do the reflection in a physical sense?

GREENHALL: It is hard to explain. If you go through the mathematics, you just find that it works. But it doesn't always work. It fails if the phony spectrum you get out of that periodic autocorrelation has some negative values. But it can be proven that, for the FD minus 1/2 and for the other model that I worked with in the paper, it always does work.

HOWE: Because ordinarily you would extend the autocovariance so that it didn't appear periodic, but rather, for example, you have truncated data – you might assume that it would exponentially decay.

GREENHALL: We are not making any assumptions. We are just trying to get an exact –

HOWE: Right. I understand that you are trying to generate flicker. Second comment: As far as the Allan deviation as a test for flicker, you commented that it was not a good test for flicker.

GREENHALL: Well, it is a necessary condition, but not sufficient.

HOWE: Right. But in long term, you know that the distribution isn't symmetric. And is the fact that the last point drops off a result of lack of number of degrees of freedom at the last point?

GREENHALL: Are you talking about IR1?

HOWE: Yes.

GREENHALL: No, it has nothing to do with that. That is just expected values. That is theory.