43rd Annual Symposium on Frequency Control - 1989

BIASES AND VARIANCES OF SEVERAL FFT SPECTRAL ESTIMATORS AS A FUNCTION OF NOISE TYPE AND NUMBER OF SAMPLES

F. L. Walls, Time and Frequency Division,
National Institute of Standards and Technology, Boulder, CO 80303

D. B. Percival, Applied Physics Laboratory,
University of Washington, Seattle, WA 98106

W. R. Irelan, Irelan Electronics,
412 Janet, Tahlequah, OK 74464

Abstract

We theoretically and experimentally investigate the biases and the variances of Fast Fourier transform (FFT) spectral estimates with different windows (data tapers) when used to analyze power-law noise types f^0, f^{-2}, f^{-3} and f^{-4}. There is a wide body of literature for white noise but virtually no investigation of biases and variances of spectral estimates for power-law noise spectra commonly seen in oscillators, amplifiers, mixers, etc. Biases (errors) in some cases exceed 30 dB. The experimental techniques introduced here permit one to analyze the performance of virtually any window for any power-law noise. This makes it possible to determine the level of a particular noise type to a specified statistical accuracy for a particular window.

I. Introduction

Fast Fourier transform (FFT) spectrum analyzers are very commonly used to estimate the spectral density of noise. These instruments often have several different windows (data tapers) available for analyzing different types of spectra. For example, in some applications spectral resolution is important; in others, the precise amplitude of a widely resolved line is important; and in still other applications, noise analysis is important. These diverse applications require different types of windows.

We theoretically and experimentally investigate the biases and variances of FFT spectral estimates with different windows when used to analyze a number of common power-law noise types. There is a wide body of literature for white noise but virtually no investigation of these effects for the types of power-law noise spectra commonly seen in oscillators, amplifiers, mixers, etc. Specifically, we present theoretical results for the biases associated with two common windows — the uniform and Hanning windows — when applied to power-law spectra varying as f^0, f^{-2} and f^{-4}. We then introduce experimental techniques for accurately determining the biases of any window and use them to evaluate the biases of three different windows for power-law spectra varying as f^0, f^{-2}, f^{-3} and f^{-4}. As an example we find with f^{-4} noise that the uniform window can have errors ranging from a few dB to over 30 dB, depending on the length of span of the f^{-4} noise.

We have also theoretically investigated the variances of FFT spectral estimates with the uniform and Hanning windows (confidence of the estimates) as a function of the power-law noise type and as a function of the amount of data. We introduce experimental techniques that make it relatively easy to independently determine the variance of the spectral estimate for virtually any window on any FFT spectrum analyzer. The variance that is realized on a particular instrument depends not only on the window but on the specific implementation in both hardware and software. We find that the variance of the spectral density estimates for white noise, f^0, is very similar for three specific windows available on one instrument and almost identical to that obtained by standard statistical analysis. The variances for spectral density estimates of f^{-4} noise are only 4% higher than that of f^0 noise for two of the windows studied. The third window — the uniform window — does not yield usable results for either f^{-3} or f^{-4} noise.

Based on this work it is now possible to determine the minimum number of samples necessary to determine the level of a particular noise type to a specified statistical accuracy as a function of the window. To our knowledge this was previously possible only for white noise — although the traditional results are generally valid for noise that varied as $f^{-\beta}$, where β was equal to or less than 4.

II. Spectrum Analyzer Basics

The spectrum analyzer which was used in the experimental work reported here is fairly typical of a number of such instruments currently available from various manufacturers. The basic measurement process generally consists of taking a string of $N_s = 1024$ digital samples of the input wave form, which we represent here by $X_1, X_2, \ldots, X_{N_s}$. The basic measurement period was 4 ms. This yields a sampling time $\Delta t = 3.90625\ \mu$s. Associated with the FFT of a time series with N_s data points, there are usually $(N_s/2) + 1 = 513$ frequencies

$$f_j = \frac{j}{N_s \Delta t}, \qquad j = 0, 1, \ldots, N_s/2.$$

The fundamental frequency f_1 is 250 Hz, and the Nyquist frequency $f_{N_s/2}$ is 128 kHz. Since the spectrum analyzer uses an anti-aliasing filter which significantly distorts the high frequency portion of the spectrum, the instrument only displays the measured spectrum for the lowest 400 nonzero frequencies, namely, $f_1 = 250$ Hz, $f_2 = 500$ Hz, …, $f_{400} = 100$ kHz.

The exact details of how the spectrum analyzer estimates the spectrum for $X_1, \ldots, X_{N_s}$ are unfortunately not provided in the documentation supplied by the manufacturer, so the following must be regarded only as a reasonable guess on our part as to its operation (see [1] for a good discussion on the basic ideas behind a spectrum analyzer; two good general references for spectral analysis are [2] and [4]). The sample mean,

$$\bar{X} \equiv \frac{1}{N_s} \sum_{t=1}^{N_s} X_t,$$

is subtracted from each of the samples, and each of these "demeaned" samples is multipled by a window h_t (sometimes called a data taper) to produce

$$X_t^{(h)} = h_t \left(X_t - \bar{X} \right).$$

The spectral estimate,

$$\hat{S}_1(f_j) = \Delta t \left| \sum_{t=1}^{N_s} X_t^{(h)} e^{-i 2\pi f_j t \Delta t} \right|^2, \qquad j = 0, 1, \ldots, N_s/2,$$

is then computed using an FFT algorithm.

The subscript "1" on $\hat{S}_1(f_j)$ indicates that this is the spectral estimate formed from the first block of N_s samples. A similar spectral estimate $\hat{S}_2(f_j)$ is then formed from the second block of contiguous data $X_{N_s+1}, X_{N_s+2}, \ldots, X_{2N_s}$. In all, there are N_b different spectral estimates from N_b contiguous blocks, and the spectrum analyzer averages these together to form

$$\hat{S}(f_j) \equiv \frac{1}{N_b} \sum_{k=1}^{N_b} \hat{S}_k(f_j). \qquad (1)$$

It is the statistical properties of $\hat{S}(f_j)$ with which we are concerned in this paper.

Unfortunately some important aspects of the windows are not provided in the documentation for the instrument. One important detail is the manner in which the window is normalized. There are two common normalizations:

$$\sum_{t=1}^{N_s} \left[h_t\left(X_t - \bar{X}\right)\right]^2 = \frac{1}{N_s} \sum_{t=1}^{N_s} \left(X_t - \bar{X}\right)^2$$

and

$$\sum_{t=1}^{N_s} h_t^2 = 1. \qquad (2)$$

The first of these is common in engineering applications because it ensures that the power in the windowed samples $X_t^{(h)}$ is the same as in the original demeaned samples; the second is equivalent to the first in expectation and is computationaly more convenient, but it can result in small discrepancies in power levels. Either normalization affects only the level of the spectral estimate and not its shape.

There are three windows built into the spectrum analyzer used here. The first is the uniform (rectangular, default) window $h_t^{(U)} = 1/\sqrt{N_s}$. The second is the Hanning data window, for which there are several slightly different definitions in the literature. In lieu of specific details, we assume the following symmetric definition:

$$h_t^{(H)} = C^{(H)}\left(1 - \cos\frac{2\pi(t-0.5)}{N_s}\right), \quad 1 \le t \le \frac{N_s}{2},$$
$$= h_{N_s-t+1}^{(H)}, \quad \frac{N_s}{2} + 1 \le t \le N_s;$$

here $C^{(H)}$ is a constant which forces the normalization in Equation (2). The third window is a proprietary "flattened peak" window, about which little specific information is available (it is evidently designed to accurately measure the heights of peaks in a spectrum).

III. Expected Value and Bias of Spectral Estimates

III.A. Theoretical Analysis

We need to assume a noise model for the X_t's in order to determine the statistical properties of $\hat{S}(f_j)$ in Equation (1). We consider three different models, each of which is represented in terms of a Gaussian white noise process ϵ_t with mean zero and variance σ_ϵ^2. The second-order properties of each model are given by a spectral density function $S(\cdot)$ defined over the interval $[-1/(2\Delta t), 1/(2\Delta t)]$ in cycles/Δt. The first model is a discrete parameter, white noise process (f^0 noise):

$$X_t = \epsilon_t \quad \text{and} \quad S(f) = \sigma_\epsilon^2 \Delta t.$$

The second model is a discrete-parameter, random-walk process (nominally f^{-2} noise):

$$X_t = \sum_{s=1}^{t} \epsilon_s \quad \text{and} \quad S(f) = \frac{\sigma_\epsilon^2 \Delta t}{4\sin^2(\pi f \Delta t)}.$$

The third model is a discrete-parameter, random-run process (nominally f^{-4} noise):

$$X_t = \sum_{r=1}^{t} \sum_{s=1}^{r} \epsilon_s \quad \text{and} \quad S(f) = \frac{\sigma_\epsilon^2 \Delta t}{16\sin^4(\pi f \Delta t)}.$$

Continuous parameter versions of these three models have been used extensively in the literature as models for noise commonly seen in oscillators.

For each of the three models we have derived expressions for $E\{\hat{S}(f)\}$, the expected value of $\hat{S}(f)$. These expressions depend on the window h_t, the number of samples N_s in each block and — in the case of a random-run process — the number of blocks N_b. The details behind these calculations will be reported elsewhere [3]; here we merely summarize our conclusions for the three models in combination with the uniform and Hanning windows and $N_s = 1024$.

First, for a white noise process,

$$E\{\hat{S}(f_j)\} = S(f_j), \qquad j = 1, 2, \ldots, 512,$$

when the uniform window is used. For the Hanning window, the above equality also holds to a very good approximation for $2 \le j \le 511$ and to within 0.8 dB for $j = 1$ and 512 (the latter is of no practical importance since the highest frequency index given by the spectrum analyzer is $j = 400$). These theoretical calculations agree with our experimental data except at f_1 (see Table 1).

Second, for a random-walk process,

$$E\{\hat{S}(f_j)\} = 2S(f_j), \qquad j = 1, 2, \ldots, 512,$$

when the uniform window is used, i.e., the expected value is *twice* what it should be at all frequencies. This theoretical result has been verified by Monte Carlo simulations, but it does *not* agree with our experimental data, which shows no significant level shift in the estimated spectrum. The source of this discrepancy is currently under investigation, but it may be due to either (a) factors in the experimental data which effectively make it band-limited, random-walk noise, i.e., its spectral shape is markedly different from f^{-2} for, say, $0 < f < f_1$ or (b) an incorrect guess on our part as to how the spectral estimate is normalized by the spectrum analyzer. For the Hanning window, we found that

$$E\{\hat{S}(f_j)\} = \begin{cases} 1.08S(f_j) & j = 1; \\ 1.48S(f_j) & j = 2; \\ 1.15S(f_j) & j = 3; \\ 1.07S(f_j) & j = 4; \\ 1.04S(f_j) & j = 5; \\ S(f_j) & 6 \le j \le 511 \text{ to within } 3\%, \end{cases}$$

i.e., $\hat{S}(f_j)$ is essentially an unbiased spectral estimate except for the lowest few frequencies. This theoretical result has been verified by Monte Carlo simulations and also agrees in general with our experimental data.

Third, for a random-run process,

$$E\{\hat{S}(f_j)\} = C_{N_b} f^{-2}, \qquad 1 \le j \le 400,$$

to a good appoximation when the uniform window is used, where C_{N_b} is a constant which depends on the number of blocks N_b and increases as N_b increases. Thus the shape of $E\{\hat{S}(f_j)\}$ follows that of a random-walk process (f^{-2}) rather than that of a random-run process (f^{-4}). This shape has been verified experimentally (see the next subsection), but the dependence of the level on N_b has not. The increase in level of $E\{\hat{S}(f_j)\}$ as N_b increases is due to the fact that the expected value of the sample variance of a block of N_s samples increases with time — by contrast, it is constant with time for the white noise and random-walk cases. For the Hanning window, we found that

$$E\{\hat{S}(f_j)\} = C'_{N_b} S(f_j), \qquad 4 \leq j \leq 400,$$

to a good approximation, where again C'_{N_b} is a constant — different from C_{N_b} — which depends on the number of blocks N_b and increases as N_b increases. For frequencies less than f_4 the theoretical results indicate significant (greater than 4%) distortion in the shape, but these do not agree in detail with the experimental values reported in Table 1. For $f_j \geq f_4$ the shape has been verified experimentally, but the dependence of the level on N_b has not. The discrepancy in level between the theoretical and experimental results is yet to be resolved, but it is probably due to a mismatch between the assumed random-run model and the true spectrum for the data (possibly band-limited random-run).

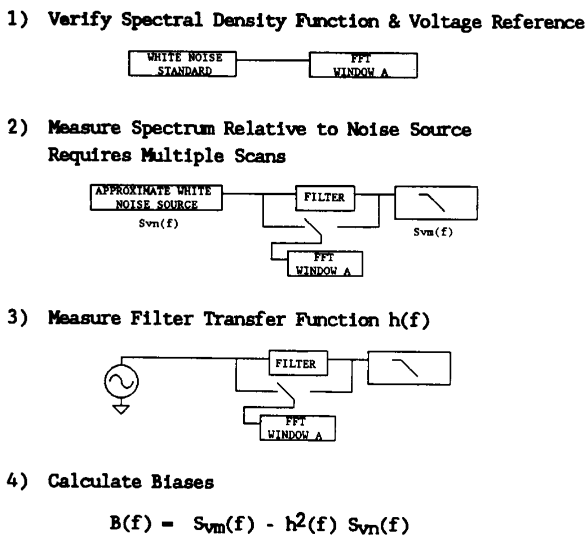

Figure 1. Outline of measurement procedure for determining the biases in spectral estimators.

III.B. Experimental Determination

The following procedure can be used to experimentally determine the bias in the spectral estimate of any noise spectrum using any window in a particular instrument. The basic concept is to implement a filter that, when applied to white noise, mimics the approximate noise spectra of interest and then measures the level of the white noise and the filter transfer function in a way which has high precision and accuracy as illustrated in Figure 1. First, the level of a known white noise is measured over a convenient range. The higher the frequency span the faster that this is accomplished. Obviously, the chosen range must be one over which the noise source is accurate. To obtain a precision of order 0.2 dB generally requires 1000 samples. This measurement verifies that the spectral density function and the internal reference voltage of the FFT are accurately calibrated and working properly. Virtually all of the windows accurately determine the value of white noise if the first few channels are ignored as explained above. Figure 2 shows the measurement of a noise source, which has been independently determined to have a noise spectral density of 99.8 dBV/Hz by the three windows. (Appendix A shows the circuit diagram for this noise source which has an accuracy of better than 0.2 dB for frequencies from 20 to 20 kHz.)

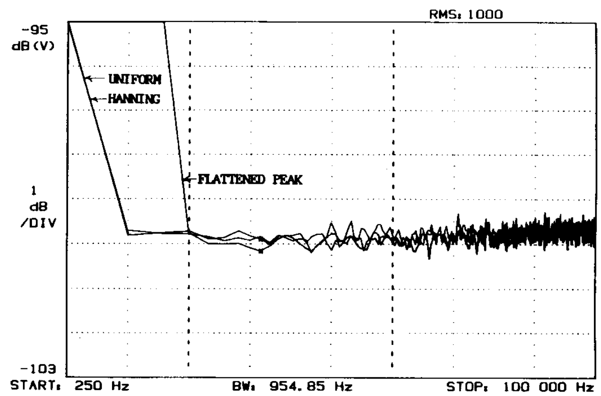

Figure 2. Spectral estimation of a white noise standard using the uniform, Hanning and the proprietary "flattened peak" windows.

Second, an approximately flat spectrum is measured over the frequency range of interest. It is not important if there are small variations in the level that change slowly over the frequency span. Third, the transfer function of the filter is determined for the frequencies of interest using a very narrow spectral source (typically an audio oscillator is sufficient). The very narrow source is accurately measured by the window since there is no problem with either high frequency or low frequency noise biasing the estimate. The use of a window with a flattened peak response is helpful but not necessary if the frequency source is sufficiently stable. This transfer function is then applied to the measured white noise spectrum in step two above. This yields a very accurate value for the "true" spectral density of the white noise source as measured through the filter. This "true value" is then compared to that obtained by the FFT analyzer. The difference between that measured in steps two and three and that measured directly with the FFT is the bias in the spectral estimate for that particular window and noise type. The accuracy of this approach comes from the fact that the calibration has been broken up into steps that can individually be determined with high precision and very small bias. The primary assumption is that the FFT analyzer is linear. Even this assumption can be checked by using precision attenuators. If the known white noise in step one does not extend to the frequencies of interest, then there is an additional assumption that the FFT is flat with frequency. This assumption is nearly always good except perhaps near the last few channels where the effect of the antialiasing filter might cause small inaccuracies.

Figures 3a and 3b show the "true" spectral estimate and the estimates as measured on a particular instrument using the uniform, Hanning, and the instrument's proprietary "flattened peak" windows for noise that varies as f^{-4} over much of the

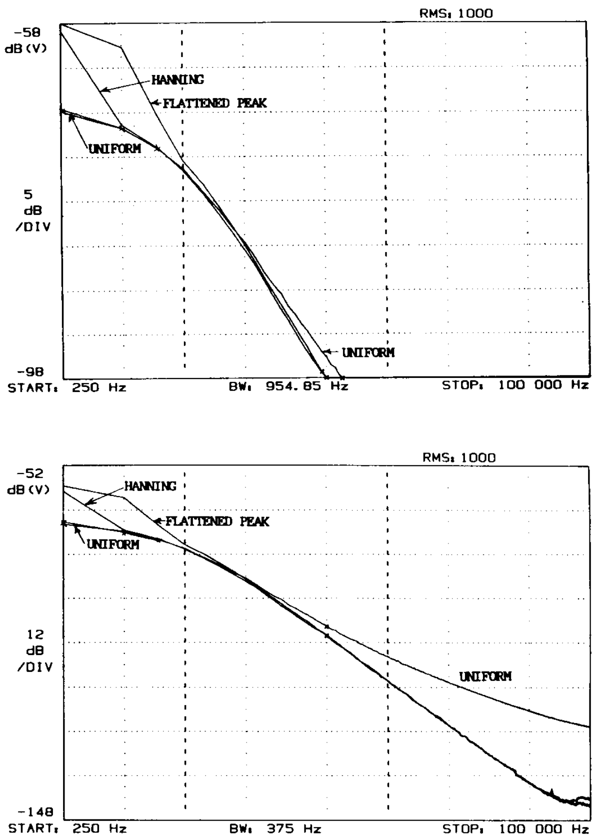

Figures 3a (top) and 3b (bottom). Difference between the true spectrum (top) which varies approximately as f^{-4} and that estimated by the uniform, Hanning and proprietary "flattened peak" windows (bottom).

range from 1 kHz to 100 kHz. The scan is 0 to 100 kHz, and 1000 samples were taken for all curves. Note the considerable difference between the spectral estimates for channels 1 to 3 for the Hanning and proprietary "flattened peak" windows. These results confirm the theoretical calculations above showing that, for the Hanning window, the first 3 channels should be ignored. For the "flattened peak" window, the first 14 channels should be ignored. For both f^{-3} and f^{-4} noise, the uniform window does not yield usable spectral estimates over *any* portion of the scan. Note in this example that at frequencies above 80 kHz there is a small step in the spectral estimates. This is due to digitizing errors of the signal due to quantization. If the digitizer had more bits, these errors would not occur. This problem of dynamic range is common whenever the spectrum of interest covers many decades. The usual solution is to use filters to divide the spectrum into various frequency range segments which are suitable for the dynamic range of the FFT.

Table 1 summarizes the measured experimental biases in the spectral estimates of a particular instrument with three different windows for power-law noise types varying from f^0 to f^{-4}. This covers most of the random types of noise found in oscillators and signal processing equipment. We do *not* advocate using the biases reported in this table to correct data — they only indicate which channels should not be relied upon for data analysis.

Table 1. Approximate Biases in FFT Spectral Estimates

noise type f^0			
channel #	uniform	Hanning	flattened peak
1	19.6 dB	19.6 dB	20.1 dB
2	small	small	16.7 dB
3	$\downarrow$	$\downarrow$	7.2 dB
4			small
5			$\downarrow$
noise type f^{-4}			
channel #	uniform	Hanning	flattened peak
1	unusable	8.6 dB	10.0 dB
2		0.4 dB	9.1 dB
3		0.4 dB	4.0 dB
4		small	1.2 dB
5		$\downarrow$	1.1 dB
6			1.1 dB
7			1.0 dB
8			0.8 dB
9			0.6 dB
10			0.6 dB
11			0.5 dB
12			0.4 dB
13			0.4 dB
14			small
15			$\downarrow$

IV. Variances of Spectral Estimates

IV.A. Theoretical Analysis

We have derived expressions for $var\{\hat{S}(f)\}$ — the variance of $\hat{S}(f)$ — for each of the three models considered in Section III.A. These expressions depend primarily on the number of blocks N_b. Again, the details behind these calculations will be reported elsewhere [3].

First, for a white noise process, the uniform window yields

$$var\{\hat{S}(f_j)\} = S^2(f_j)/N_b, \qquad 1 \le j \le 511,$$

while the Hanning window yields

$$var\{\hat{S}(f_j)\} = \begin{cases} 0.69 S^2(f_j)/N_b, & j = 1; \\ S^2(f_j)/N_b, & 2 \le j \le 510; \\ 1.03 S^2(f_j)/N_b, & j = 511. \end{cases}$$

These results are consistent with our experimental results and with standard statistical theory.

Second, for a random-walk process, the uniform window yields

$$var\{\hat{S}(f_j)\} = 5 S^2(f_j)/N_b, \quad 1 \le j \le 511,$$

while the Hanning window yields

$$var\{\hat{S}(f_j)\} = \begin{cases} 1.30S^2(f_j)/N_b, & j=1; \\ 2.20S^2(f_j)/N_b, & j=2; \\ 1.31S^2(f_j)/N_b, & j=3; \\ 1.15S^2(f_j)/N_b, & j=4; \\ 1.09S^2(f_j)/N_b, & j=5; \\ 1.06S^2(f_j)/N_b, & j=6; \\ 1.04S^2(f_j)/N_b, & j=7; \\ S^2(f_j)/N_b, & 8 \le j \le 511 \text{ to within } 3\%. \end{cases}$$

Except for the few lowest frequencies, the results for the Hanning window agree with our experimental results and with standard statistical theory; however, the factor of five in the variance for the uniform window disagrees with our experiments and with standard theory (although it has been verified by Monte Carlo techniques). The cause of this discrepancy is under investigation, but we think it is due to the band-limited nature of the experimental data.

Third, for a random-run process, the variance computations are not useful since the variance is dominated by the fact that the expected value of the sample variance for each block of samples increases with time. The agreement which we found between standard statistical theory and our experimental results on the $1/N_b$ rate of decrease of variance is undoubtedly due to the band-limited nature of the experimental data. We will attempt to verify these conclusions in the future using Monte Carlo techniques.

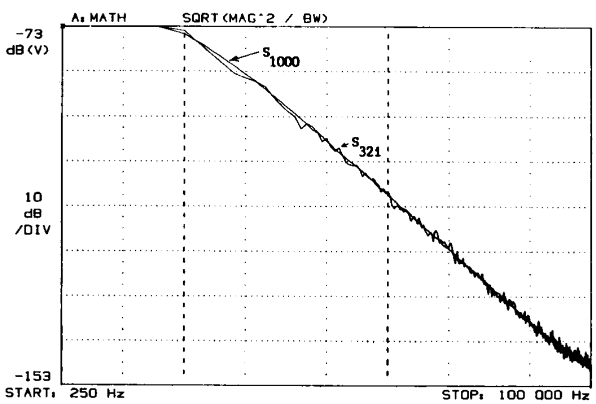

Figure 4. Comparison of the spectral estimate of f^{-4} power-law noise with 1000 samples with that obtained with 32 samples. The text explains how these two curves are used to obtain the fractional RMS confidence of the spectral estimate for 32 samples.

IV.B. Experimental Determination

The following procedure can be used to experimentally determine the variance of the spectral estimates of virtually any type of noise spectrum with any type of window for a particular instrument. Since the spectral density of interest is in general nonwhite, we must determine both the "true value" and a way to normalize the fractional error of the estimate as a function of the number of samples. This can be done by making use of the above theoretical analysis that shows that the variance should decrease as the square root of the number of samples since they are approximately statistically independent (in fact, exactly so in the cases of white and random-walk noise). As an example we have chosen to take $N_b = 1000$ blocks of the various power-law noise types examined in III.B above and compare the value of the spectral estimate with that obtained from $N_b = 32$ blocks (see Figure 4). Since the variance of the 1000 block data is about 32 times smaller than that of the 32 block data, it can serve as an accurate estimate of the "true value." Let $\hat{S}_{1000}(f_j)$ represent this quantity at the j-th channel (frequency). By subtracting the 1000 block data from the 32 block data at the j-th channel, we then have one estimate of the error for the 32 block data; by repeating this procedure over N_c different channels and N_r different replications, we can obtain accurate estimates of the variance for the 32 block data. Let $\hat{S}_{32i}(f_j)$ represent the spectral estimate for the 32 block data at the j-th channel and the i-th replication. To compensate for the variation in the level of the spectral estimates with channel, it is necessary to divide the error at the j-th channel by the "true value" $\hat{S}_{1000}(f_j)$. The mean square fractional error of the 32 block data for the noise type under study is given by

$$\sigma_{32}^2 = \frac{1}{N_r N_c}\sum_{i=1}^{N_r}\sum_{j}\left(\frac{\hat{S}_{32i}(f_j)-\hat{S}_{1000}(f_j)}{\hat{S}_{1000}(f_j)}\right)^2 \approx \frac{var\{\hat{S}_{32i}(f_j)\}}{S^2(f_j)}.$$

It is assumed that all channels with bias — as indicated in Table 1 — have been excluded in the summation over j. It is also important that the changes in the spectral density not exceed the dynamic range of the digitizer because under this condition the quantization errors — in addition to causing biases in the spectral estimates as discussed earlier — can lead to situations where the variance does not improve as N_b increases. These values can be scaled to any number of blocks N_b if care is taken to avoid these quantization errors. Upper and lower approximate 67% confidence limits for $S(f_j)$ — the true spectral density at channel j — using Hanning, uniform and the proprietary "flattened peak" windows for N_b approximately independent blocks are given by

$$\hat{S}(f_j)\,(1 \pm V(\alpha, N_b))$$

where $\hat{S}(f_j)$ is the spectral estimate given by Equation (1) and $V(\alpha, N_b)$ is the fractional variance given in Table 2 for f^α and α = 0, -2, -3 and -4 (these results were obtained by averaging over $N_r N_c$ = 1200 channels). The variances obtained are very close to those obtained from standard statistical analysis for white noise, i.e.,

$$\hat{S}(f_j)\left(1 \pm \frac{1}{\sqrt{N_b}}\right).$$

Table 2. Confidence Intervals for FFT Spectral Estimates

power law noise type	uniform	window Hanning	flattened peak
f^0	$1.02/\sqrt{N_b}$	$0.98/\sqrt{N_b}$	$0.98/\sqrt{N_b}$
f^{-2}	$1.02/\sqrt{N_b}$	$1.04/\sqrt{N_b}$	$1.04/\sqrt{N_b}$
f^{-3}	unusable	$1.04/\sqrt{N_b}$	$1.04/\sqrt{N_b}$
f^{-4}	unusable	$1.04/\sqrt{N_b}$	$1.04/\sqrt{N_b}$

V. Conclusions

We have introduced experimental techniques to evaluate the statistical properties of FFT spectral estimates for common noise types found in oscillators, amplifiers, mixers and similar devices, and we have compared these with theoretical calcuations. We have used these techniques to study the biases and variances of FFT spectral estimates using the uniform, Hanning, and a proprietary "flattened peak" window. The theoretical analysis was greatly hampered because the instrument manufacturer does not disclose the exact form of the "flattened peak" window or the normalization procedure for the other windows. Nevertheless, we obtained fair agreement between the theoretical and the experimental analysis. The variances of the spectral estimation were virtually identical to a few percent for f^0 to f^{-4} noise except for the uniform window which is incapable of measuring noise which falls off faster than f^{-2}. There was a very large difference in the biases of the first few channels for the three windows. The Hanning window showed significant biases in the first 3 channels while the proprietary "flattened peak" window showed large biases for f^{-4} noise even up to channel 13. The Hanning window therefore yields useful information over three times wider frequency range than the proprietary "flattened peak" window. In the particular instrument studied, the proprietary "flattened peak" window is the best choice for estimating the height of a narrow band source, while the Hanning window is by far the best choice for spectral analysis of common noise types found in oscillators, amplifiers, mixers, etc. We have also shown that the 67% confidence levels for spectral estimation as a function of the number of contiguous nonoverlapping blocks, N_b, is approximately given by

$$S = S_m \left(1 \pm \frac{0.98}{\sqrt{N_b}}\right)$$

for white noise (f^0) and by

$$S = S_m \left(1 \pm \frac{1.04}{\sqrt{N_b}}\right)$$

for noise types f^{-2} to f^{-4}. This agrees to within 4% of that found by standard statistical analysis for white noise. Using this data one can now determine the number of samples necessary to estimate — to a given level of statistical uncertainty — the spectrum of the various noise types commonly found in oscillators, amplifiers, mixers, etc.

References

[1] D. Babitch and J. Oliverio, "Phase Noise of Various Oscillators at Very Low Fourier Frequencies," in *Proceedings of the 28th Annual Symposium on Frequency Control*, 1974, pp. 150–159.

[2] A. V. Oppenheim and R. W. Schafer, *Discrete-Time Signal Processing.* Englewood Cliffs, New Jersey: Prentice Hall, 1989, ch. 11, pp. 695–767.

[3] D. B. Percival and F. L. Walls, "Spectral Analysis of Power-Law Processes with a Spectrum Analyzer," manuscript in preparation, 1989.

[4] M. B. Priestley, *Spectral Analysis and Time Series.* London: Academic Press, 1981, ch. 7, pp. 502–612.

Appendix. Precision Noise Source

Figure 5 shows the circuit diagram of a precision noise source whose spectral density can be determined from first principles to ±0.2 dB over the frequency range from 20 Hz to 20 kHz. The spectral density is basically given by the Johnson noise of the 10^5 ohm resistor, $V_n^2 = 4kTR$, where T is in Kelvin, and k is Boltzmann's constant. Corrections due to the input noise voltage and noise current of the amplifier amount to about 0.2 dB for the circuit elements shown. All resistors are precision 1% metal film resistors. The output level can be switched from -100 dBV/Hz to -80 dBV/Hz. By adjusting the noise-gain capacitors one can make the noise spectrum flat to within 0.3 dB out to 200 kHz. There is also provision to measure the input noise voltage of the amplifier by shorting the input to ground or the combined noise voltage and noise current by switching a 220 pF capacitor into the input instead of the 10^5 ohm noise resistor.

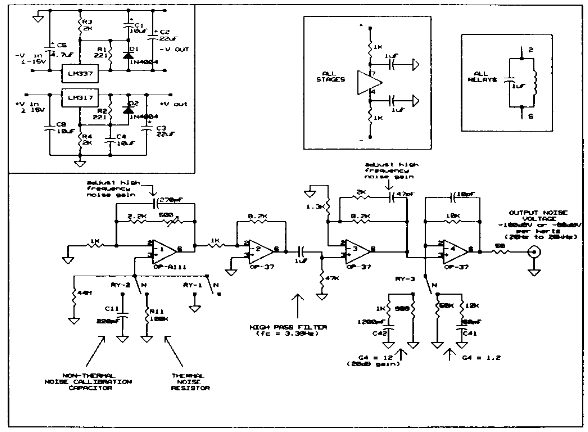

Figure 5. Circuit diagram of a precision noise source.