DATA WINDOWING AND THE FX CORRELATOR

FRED SCHWAB

February 4, 1987

Introduction. Here I want to discuss data tapering, or windowing, as it relates to the "FX" spectral processor design under consideration for the VLBA. Since sidelobes of sharp spectral features contaminate adjacent regions of the cross-spectral estimates, control of these sidelobes—which is most easily accomplished via time-domain windowing—may be one of the critical factors in achieving high spectral dynamic-range.

A summary of the current design for the FX processor is given in Escoffier and Greenberg's VLBA Memorandum No. 72 (July, 1986). In this design, pairs of quantized recorded signals are read from magnetic tape, with appropriate synchronization and delay. "Fringe rotation" is applied, and re-quantized samples of the data, representing finite length segments of two (complex) random processes, $x(t)$, $y(t)$, $0 \leq t \leq T$, are fed in chunks, K of them, each of length L, to the FFT processor. The estimate of the cross-spectrum $G_{xy}(f)$ is an average of periodogram estimates based on these (possibly overlapping) data chunks,

$$\widehat{G}_{xy}(f) = \frac{1}{K} \sum_{k=1}^{K} X_k(f) Y_k^*(f),$$

where

$$X_k(f) = \int_0^L w(t - L/2)\, x(t + (k-1)b)\, e^{-2\pi i f t}\, dt,$$

and

$$Y_k(f) = \int_0^L w(t - L/2)\, y(t + (k-1)b)\, e^{-2\pi i f t}\, dt.$$

Here, $w(t)$ denotes the data window, or taper; and b is the time offset between adjacent data chunks. $(1 - b/L)$ is the fractional overlap: $b = L$ corresponds to 0% overlap and $b = L/2$ to 50% overlap (these are probably the only two choices that will be allowed by the VLBA processor; further, some modes of operation may offer no capability for overlap).

In the above I have given a continuous formulation, but we are actually dealing with a set of discrete samples, $x(t_n)$, $y(t_n)$, $n = 0, \ldots, N$, uniformly spaced in time, and our Fourier integrals above actually correspond to FFT's in the FX processor. This is unimportant in most of the discussion below.

To see why data windows other than $w(t) \equiv$ constant might be appropriate, consider the case $K = 1$. In this case, the cross-spectral estimate $\widehat{G}_{xy}(f)$ is simply a raw cross-periodogram based on N data samples—not an average of periodograms. According to the statistical literature (e.g., [6, p. 211], [7, p. 265 ff.]), the raw (i.e., untapered) cross-periodogram is not a consistent estimator of $G_{xy}(f)$: as you allow N to increase, the variance of the estimate does not decrease. (Rather, what happens is that the additional data allow you to obtain essentially statistically independent estimates of G at a larger number of frequencies than you could have before.)

There are a number of ways to reduce the variance. One is to average the cross-spectral estimates over frequency. Since convolution is a form of averaging—transforming to the lag domain, tapering the cross-correlation estimate, and transforming back, is a special case of averaging. Another method which is essentially equivalent to averaging in the spectral domain is to apply a time domain taper before computing the periodogram; what this yields is an estimate of the convolution of $G_{xy}(f)$ with the square of the Fourier transform of the time domain taper, i.e., with $|W(f)|^2$; this is essentially equivalent to tapering the cross-correlation function with the autocorrelation of the time domain window. Another way is to divide the data into $K > 1$ possibly overlapping chunks, compute K

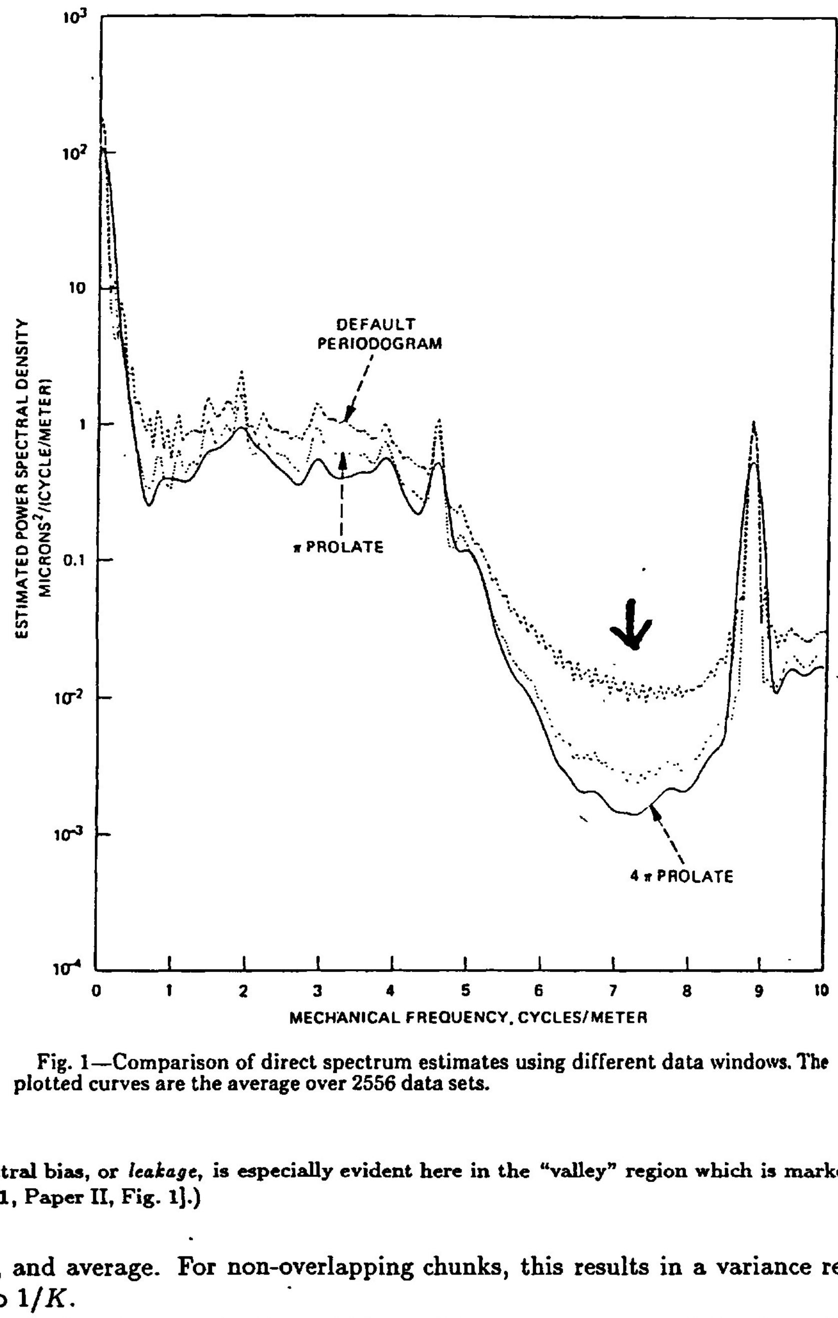

Fig. 1—Comparison of direct spectrum estimates using different data windows. The plotted curves are the average over 2556 data sets.

Figure 1. Spectral bias, or *leakage*, is especially evident here in the "valley" region which is marked by an arrow. (Adapted from [1, Paper II, Fig. 1].)

periodograms, and average. For non-overlapping chunks, this results in a variance reduction by a factor equal to $1/K$.

The method of variance reduction which would seem most natural for the FX correlator is a combination of periodogram averaging with time domain tapering. Since the expected value of a cross-periodogram estimate is the true cross-spectrum convolved with the spectral window $|W(f)|^2$ (see [1]), the sidelobes of the spectral window are a source of bias in the neighborhood of any sharp spectral feature. At the expense of spectral resolution, one can simultaneously achieve variance reduction *and* bias reduction, via appropriate segmentation, windowing, and overlap. (In the literature the method we are using is often referred to as the WOSA method: the method of Weighted Overlapped Segment Averaging. It's also sometimes called the Welch Method, having first been described at length by Welch [8] in 1967.) When K is large—as it will be, I believe, in all VLBA applications—time domain windowing is used more for purposes of bias reduction than for variance reduction.

For an illustration of this spectral bias (which is often called *spectral leakage* in the literature), see Figures 1 and 2. The first figure shows (auto-) spectral estimates computed using different data windows. The curve labeled "default periodogram" corresponds to a uniform time domain window—

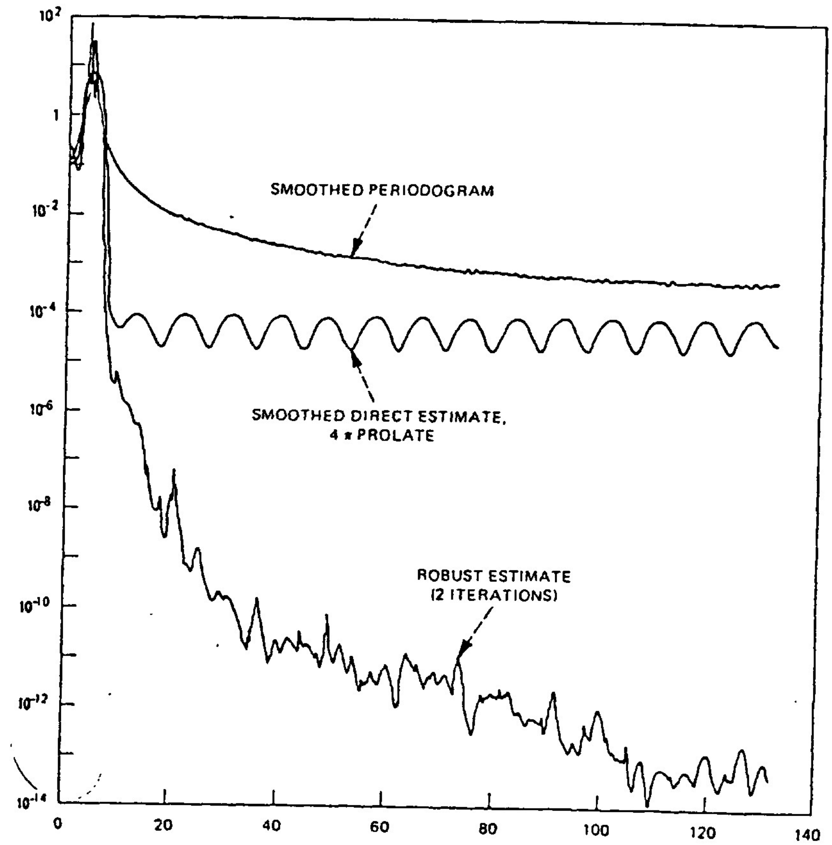

Fig. 18—Comparison of three estimates of spectrum for the data shown in Fig. 12. The systematic oscillations in the center curve are a result of interactions between the two outliers.

Figure 2. An extreme example of the problem of spectral bias. A sophisticated data window was used to obtain the spectral estimate shown by the lower curve. Almost all of the spectral features are obscured in the upper curve, which is a smoothed version of the average of "raw" periodograms (no tapering). (Adapted from [1, Paper II, Fig. 18].)

i.e., to no tapering; the other curves correspond to heavier tapering. The default periodogram estimates in the "valley" which I have marked by an arrow are biased heavily upward by the relatively high sidelobes of the (sinc^2) spectral window, because of the sharp spectral peaks surrounding the valley. An extreme example of this effect is shown in the second figure.

Also evident from Figure 2 is that workers in other fields are interested in obtaining much higher spectral dynamic range than we probably are (note that this spectrum, which is plotted on a logarithmic scale, covers nearly 16 decades in power!). Thus, many of the spectral windows that are described in the literature (probably including Hanning) should not be of much interest to us: they sacrifice too much in resolution, in order to achieve extremely high rates of sidelobe decay. Nevertheless, D. J. Thompson in [1] comments that

> Very few spectra resulting from physical processes are so uninteresting that the "elimination" of tapering is ever advisable.

Quantized windows. The currently proposed design for the FX processor uses crude precision binary floating-point arithmetic within the FFT modules: 5 bit mantissas and 4 bit exponents (counting the sign bits) for data and for intermediate results, with the real and imaginary parts of each complex datum sharing the same exponent. It uses 6 bit mantissas and 4 bit exponents to represent the FFT "twiddle factors" (roots of unity). Data sequences at the input to the FFT modules can be tapered by a quantized data window represented in floating-point arithmetic with 5 bit mantissas and 4 bit exponents. (Some numerical simulations performed by John Benson have suggested that 8 bit mantissas may be required.)

One question is whether the standard data windows, represented with this degree of precision, retain their gross properties and their overall effectiveness. To answer it, I quantized a few of the standard data windows, computed the corresponding spectral windows (the squared magnitude of their Fourier transforms), and compared these with the spectral windows corresponding to the unquantized data windows.

A sample result is shown in Figure 3. The unquantized data window is shown at top left, and the quantized version at top right. The middle two plots show the inner portions of the corresponding spectral windows. I've labeled the abscissae of the plots of the spectral windows in units of $1/T$, where T (*for the next few paragraphs only—I don't want to relabel my plots*) represents the total width, in time, of the data window. (With factor-of-2 zero-padding, the output points of the FX processor would be spaced at (the customary) increments of $\frac{1}{2T}$; without zero-padding the spacing would be $1/T$.)[1] The bottom pair of plots shows the far sidelobes of the spectral windows: what's interesting to note here is the rate of sidelobe decay and the degree of "hash" in the spectral window corresponding to the quantized data window.

In the Appendix I show a number of other examples. In all cases, the first few sidelobes of the spectral windows show very little distortion due to quantization; in particular, the first two sidelobes generally agree very closely. By a frequency of $5/T$ considerable distortion begins to show up. For each window pair the spectral envelopes of the far sidelobes generally are fairly similar, though the envelope of the quantized data window generally is more reluctant to fall below about 10^{-6}.

For comparison, I show in Figure 4 the (sinc^2) spectral window which corresponds to uniform weighting. Its first sidelobe is at a level of ≈ 0.0456 (-13.41 dB). At a frequency increment $\approx 10/T$ from a strong spectral line the sidelobes would limit the achievable dynamic range to 1000 : 1, roughly speaking.

Quantization of typical data windows to 5 bit mantissa, 4 bit exponent, appears to do little harm except to the far sidelobes of the spectral windows. Since it is easy to keep these below, say, 10^{-5} or 10^{-6}, and since we're probably content with spectral dynamic range limits of 10^4 or so, this quantization evidently is not a problem. Many of the typical data windows (Hanning, for example) have far higher rates of sidelobe decay than we would require, at considerable expense in spectral resolution (see page A–1 of the Appendix). For example, the window shown in Figure 3 has a first sidelobe level that approximately matches that of the Hanning window, but its main lobe is a bit narrower.

Often, in other applications, it doesn't matter very much if the resolution is decreased a bit due to tapering, because one can then just use longer FFT's, but fewer of them, to recover the lost resolution. In the case of the FX processor, however, the FFT size is restricted (tentatively, at least) to 2048 points.

Crude-precision FFT of a quantized window. I hadn't intended in this memo to investigate the implications of crude-precision arithmetic in the FFT modules of the FX processor, since that's a matter distinct from windowing considerations. However, a plot that I showed off at meetings of the VLBA correlator group generated a bit of interest at those meetings, so I'm including it here in Figure 5. Figure 5 compares a high-precision FFT (essentially an exact Fourier transform) of a quantized data window with a crude-precision (5 bit mantissa) FFT of the same window. The crude-precision FFT was computed using a subroutine from John Benson's simulation program on the Convex. What is of interest is the "ringing" artifacts evident in the lower right hand panel of the Figure.

---

[1] To generate the plots shown in Figure 3 and in the Appendix, I used 2048-point discrete data windows and 16384-point FFT's. This yielded points spaced at an increment of $\frac{1}{8T}$. Factor-of-2 zero-padding is customary because it leads to no information loss. The matter of zero-padding—which we have first seen discussed by O'Sullivan (NFRA Note No. 375) and then in VLBA Correlator Memoranda Nos. 66 and 67, where it is dubbed "O'Sullivan zero-padding"—too is discussed in detail in a 1967 paper, [9], by Bingham *et al.*

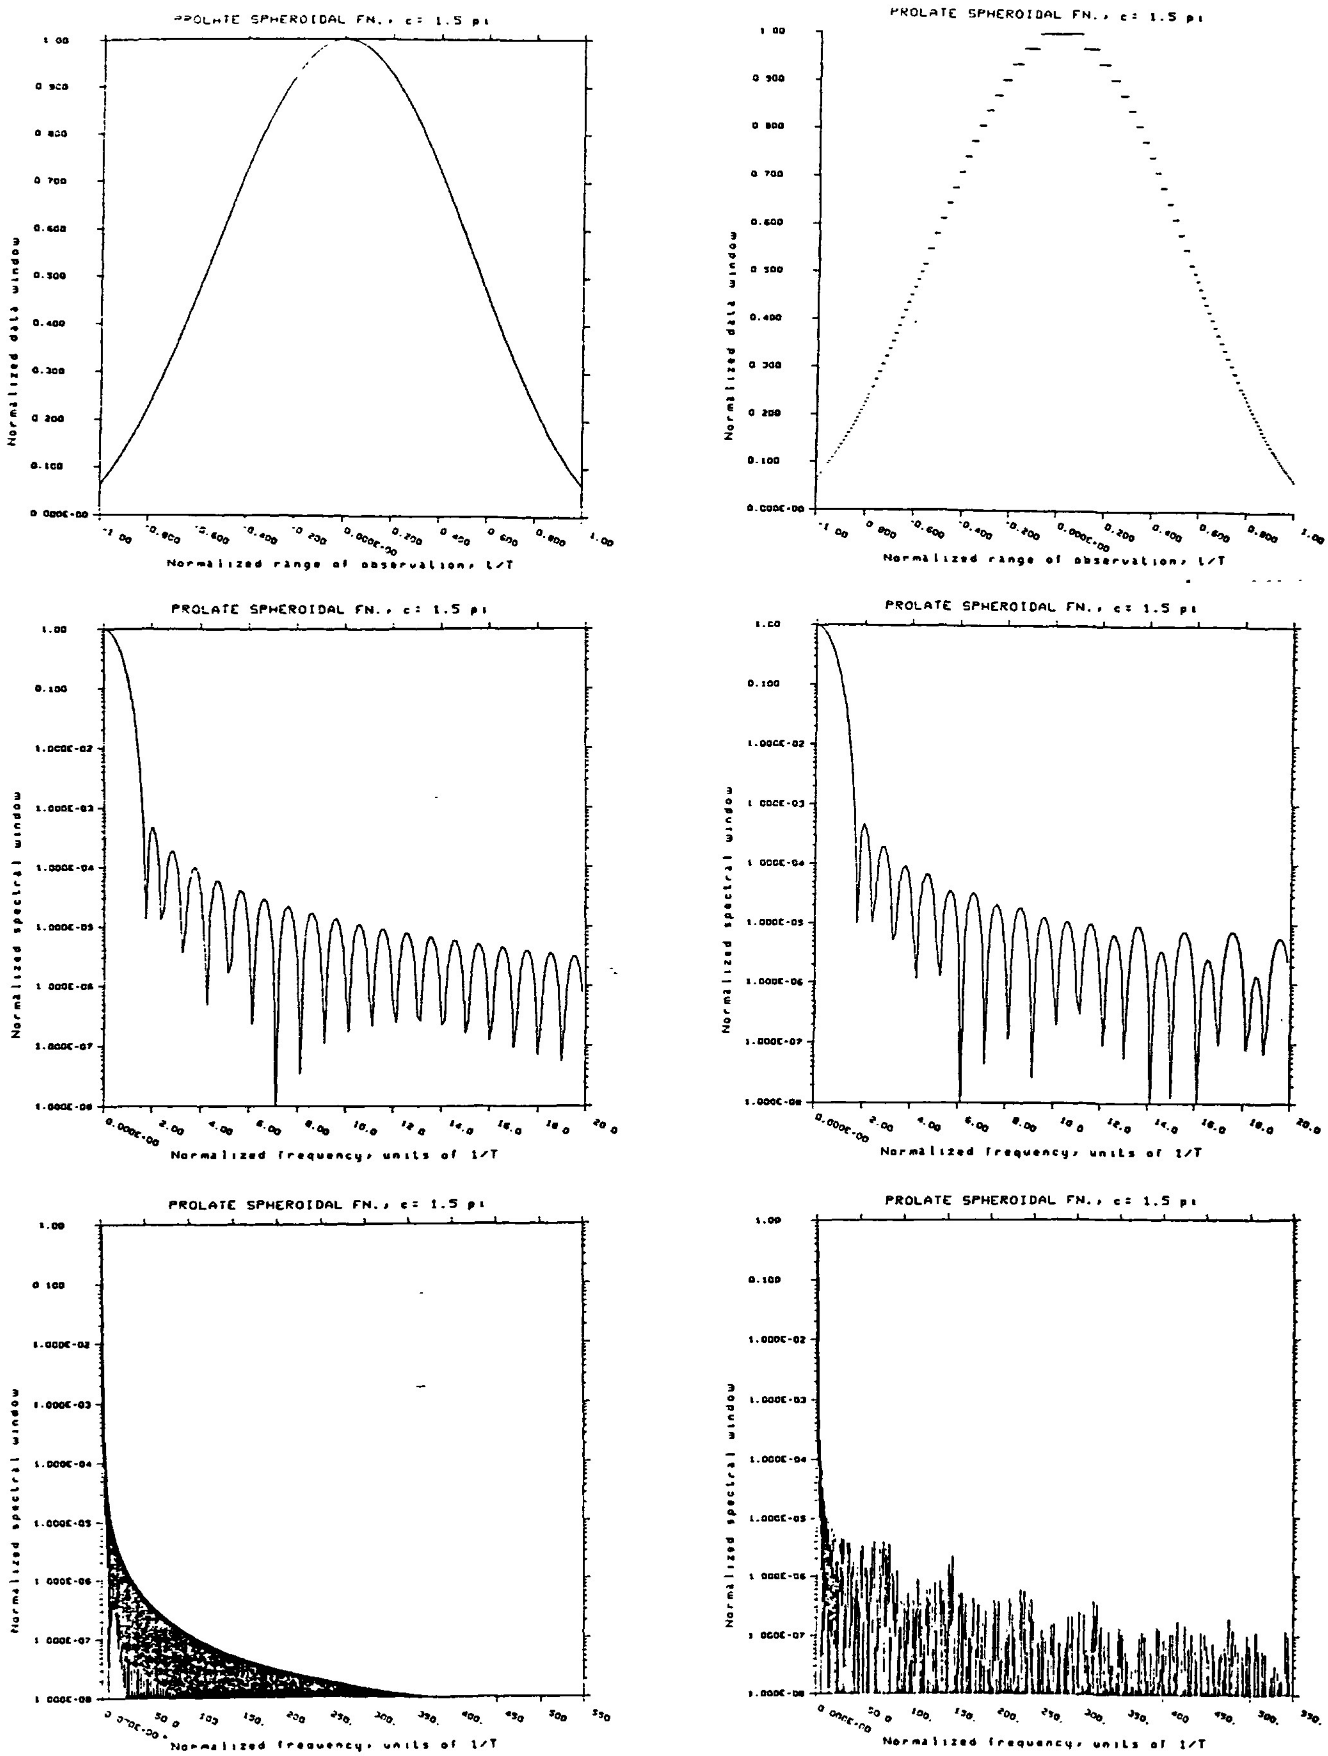

Figure 3. A typical window. The unquantized data window and the quantized version thereof are shown at top left and top right, respectively, and the corresponding spectral windows are shown underneath. Near sidelobes are shown in the middle pair of plots, and far sidelobes at the bottom.

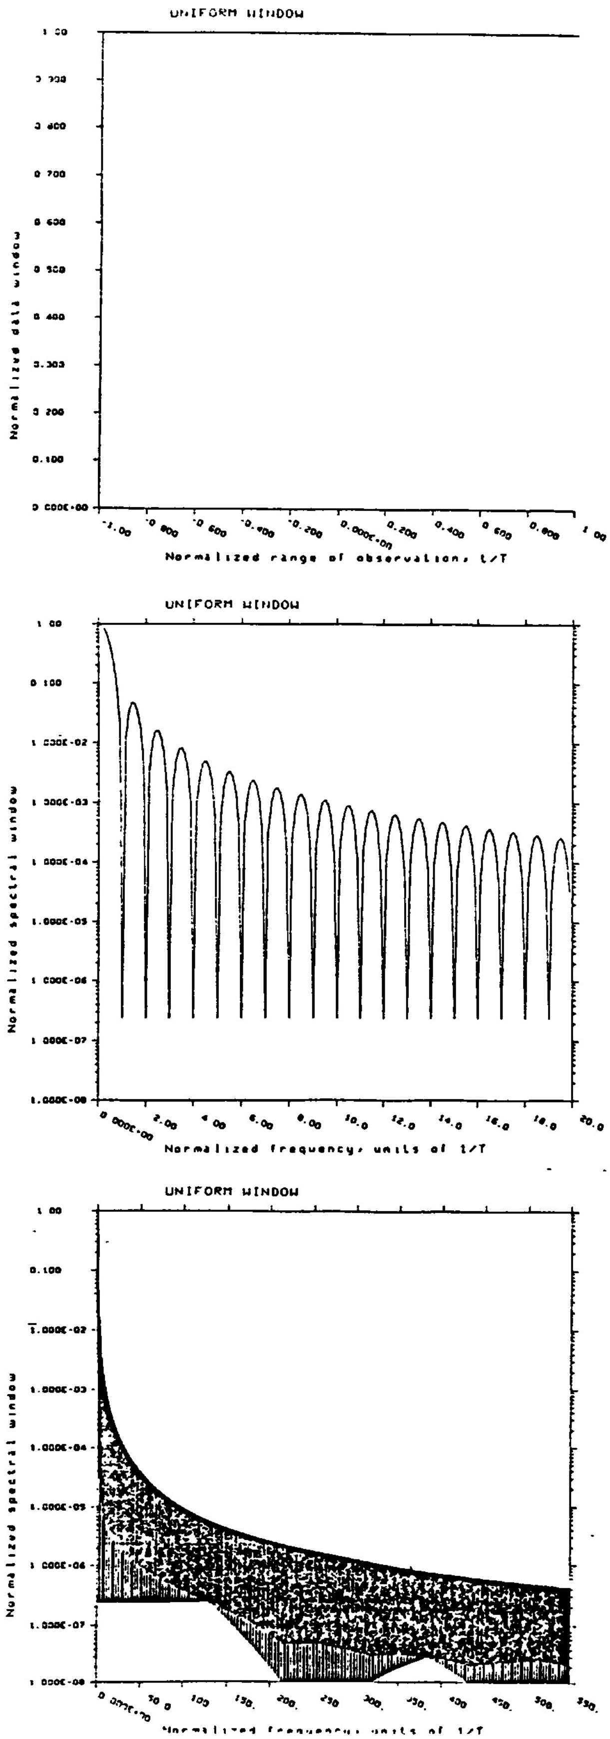

Figure 4. A uniform data window and the corresponding sinc^2 spectral window.

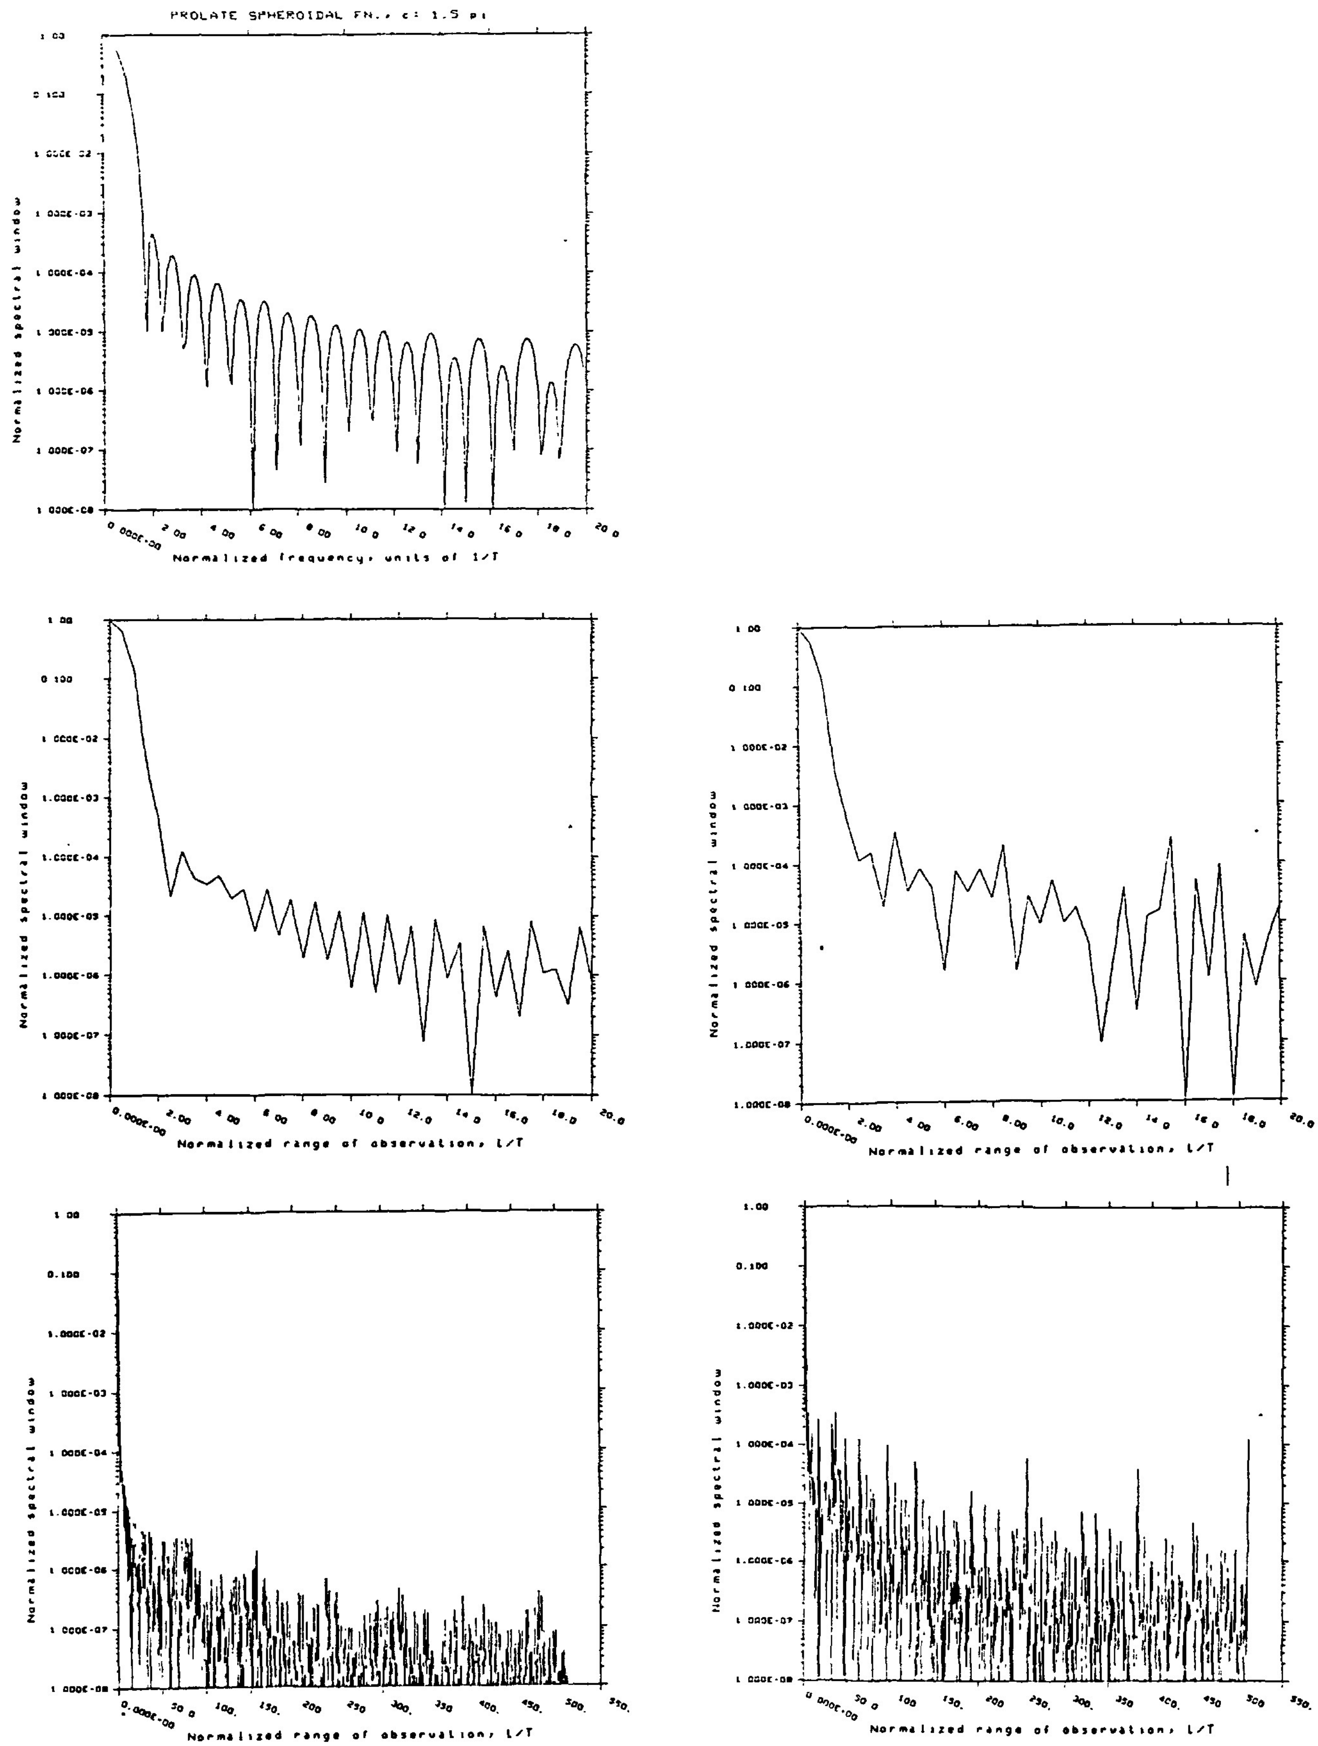

Figure 5. *(Left Panel)* High-precision floating point FFT's of a 1024-point quantized data window (16384-point transform at top, 2048-point transforms at middle and bottom—for comparison with right panel). *(Right Panel)* Crude precision (5 bit mantissa) 2048-point floating point FFT of a 1024-point quantized data window. The data window is the same as that used in Figure 3.

Overlap. Nuttall [2–4] and Thompson [1] present a number of useful results pertaining to the amount of overlap that is required in order to most efficiently estimate $G_{xy}(f)$. I want in this section to summarize some of their results. Further information is available in the published literature, [1] and [5]. Appendix B contains abstracts of the three technical reports, [2–4]; as these reports are not readily available, I can provide copies on request.

Let W denote the Fourier transform of the data window w, which we'll assume to be real-valued and even. The mean of the cross-spectral estimate $\widehat{G}_{xy}(f)$ is given by the convolution of G_{xy} with the spectral window $|W(f)|^2$,

$$\mathcal{E}\left\{\widehat{G}_{xy}(f)\right\} \simeq \int_{-\infty}^{\infty} G_{xy}(\mu)|W(f-\mu)|^2\,d\mu\,.$$

Then assuming that the width B, say the half-power width, of $|W|^2$ is narrower than the finest detail in the cross-spectrum (B is of order $1/L$) and that w has been normalized so that $\int_{-\infty}^{\infty}|W(f)|^2\,df = 1$,

$$\mathcal{E}\left\{\widehat{G}_{xy}(f)\right\} \simeq G_{xy}(f)\int_{-\infty}^{\infty}|W(\mu)|^2\,d\mu = G_{xy}(f)\,.$$

For the random variable $\widehat{g}(f) = \widehat{G}_{xy}(f) - \mathcal{E}\left\{\widehat{G}_{xy}(f)\right\} \simeq \widehat{G}_{xy}(f) - G_{xy}(f)$, one has

$$\mathcal{E}\left\{|\widehat{g}(f)|^2\right\} \simeq G_{xx}(f)G_{yy}(f)\left(\frac{1}{K}+\frac{2}{K}\sum_{k=1}^{K-1}\left(1-\frac{k}{K}\right)L_w^2(kb)\right), \tag{1}$$

and

$$\mathcal{E}\left\{\widehat{g}^2(f)\right\} \simeq G_{xy}^2(f)\left(\frac{1}{K}+\frac{2}{K}\sum_{k=1}^{K-1}\left(1-\frac{k}{K}\right)L_w^2(kb)\right). \tag{2}$$

Here $L_w \equiv w * w$ denotes the equivalent lag-domain window, the autocorrelation of the data window (which is also the self-convolution of the window, since w is even). Nuttall's derivation of (1) and (2) requires three assumptions: that $x(t)$ and $y(t)$ are jointly Gaussian; that the frequency f of interest is not too close to 0, i.e., that it exceeds the resolution bandwidth B; and that width B is narrower than the sharpest detail in each of G_{xy}, G_{xx}, and G_{yy}.

Using the above results one can characterize the stability of $\widehat{g}(f)$ in terms of an equivalent number of degrees of freedom, given by twice the squared magnitude of the expectation, divided by the variance,

$$\text{EDF} = \frac{2|\mathcal{E}\{\widehat{g}\}|^2}{\mathcal{E}\left\{|\widehat{g}-\mathcal{E}\{\widehat{g}\}|^2\right\}}.$$

The number of equivalent degrees of freedom can be written as $\text{EDF} = |\gamma_{xy}(f)|^2\mathcal{K}$, where

$$\gamma_{xy}(f) \equiv \frac{G_{xy}(f)}{\sqrt{G_{xx}(f)G_{yy}(f)}} \qquad \text{and} \qquad \mathcal{K} \equiv \frac{2K}{1+2\sum_{k=1}^{K-1}\left(1-\frac{k}{K}\right)L_w^2(kb)}.$$

The variance of the amplitude of the cross-spectral estimate is given (approximately) by

$$\mathcal{V}\left\{\left|\widehat{G}_{xy}(f)\right|\right\} \simeq G_{xx}(f)G_{yy}(f)\frac{1+|\gamma_{xy}(f)|^2}{\mathcal{K}},$$

and the variance of the phase estimate by

$$\mathcal{V}\left\{\arg\widehat{G}_{xy}(f)\right\} \simeq \frac{1-|\gamma_{xy}(f)|^2}{|\gamma_{xy}(f)|^2\mathcal{K}}.$$

$\mathcal{K}$ depends on the number K of data segments, the time offset b between segments, and the autocorrelation L_w of the data window. Since the variances depend inversely on $\mathcal{K}$, the idea now is to maximize K while avoiding excess computational effort. One might ask: "For fixed K and a given choice of data window, what is the maximum attainable number $\mathrm{EDF}_{\max}$ of equivalent degrees of freedom?", "For fixed K and a given choice of data window, what amount of overlap would suffice to realize, say, 99% or 95% of $\mathrm{EDF}_{\max}$?", "Ditto, in the limit $K \to \infty$?", or "For a given choice of data window and a fixed amount of overlap, say, 50%, what fraction of $\mathrm{EDF}_{\max}$ is attained?".

Nuttall presents numerous tables which bear on these and related questions. I'll summarize some of his results:

(1) He shows that, except for small time-bandwidth products BT—where, as above, B refers to the resolution half-power bandwidth and T to the total duration of the data—the same EDF is attainable with the "FX" approach, via appropriate overlap, as with the "XF".

(2) All typical data windows have essentially the *same* variance reduction capability when compared under the same frequency resolution constraints. An approximate rule-of-thumb is $\mathrm{EDF}_{\max} \simeq 3(BT - 1)$.

(3) For typical data windows and $BT \gtrsim 4$, the fractional overlap that is required in order to attain a specified fraction of $\mathrm{EDF}_{\max}$ is essentially constant (as a function of BT).

(4) For large BT the number K of data segments required to realize $\mathrm{EDF}_{\max}$ is significantly greater than the number required to realize, say, 95% or 99% of $\mathrm{EDF}_{\max}$. Typically, $K \approx 1.75BT$ segments are required to attain 99% of $\mathrm{EDF}_{\max}$.

(5) A perhaps counter-intuitive result: For finite BT, as the amount of overlap approaches 100%, EDF reaches a maximum and then decreases slightly. This effect is not evident in Figure 6 below, because the plots shown there correspond to the limiting case $BT \to \infty$.

(6) Nuttall's early report [2], the source of his tabular data, does not present data for spectral windows of much practical interest, apart from the Hanning window. For that case he shows that the required fractional overlap for 99% of $\mathrm{EDF}_{\max}$ is 61%, that 92% of $\mathrm{EDF}_{\max}$ is attainable with 50% overlap, and that $\sim$ 100% of $\mathrm{EDF}_{\max}$ is attainable with 62.5% overlap. Since we would require, in our application, much less extreme tapering than Hanning, I would guess that $\sim$ 95% or more of $\mathrm{EDF}_{\max}$ would be attainable with 50% overlap. $0.95\mathrm{EDF}_{\max}$ corresponds to a signal-to-noise loss of 2.6%.

D. J. Thompson in [1] presents another way of measuring the effectiveness of overlap and windowing, and his paper includes a particularly useful set of plots, which I will duplicate here. From above, the variance (as a function of the amount of overlap) of a spectral estimate based on K subsets of data is proportional to

$$V_K(b) = \frac{1}{K} + \frac{2}{K} \sum_{k=1}^{K-1} \left(1 - \frac{k}{K}\right) L_w^2(kb).$$

If enough data now are added to enable one to calculate $K + 1$ subsets, a measure of the relative gain in information is given by

$$\Delta I_K(b) = \frac{1}{b} \left(\frac{1}{V_{K+1}(b)} - \frac{1}{V_K(b)} \right).$$

As $K \to \infty$, ΔI_K approaches

$$\Delta I_\infty(b) = \frac{1}{b} \frac{1}{1 + 2\sum_{k=1}^{[T/b]} L_w^2(kb)}.$$

Plots of ΔI_∞, corresponding to a few typical data windows, are shown in Figure 6. The point is that one should be far enough to the left on the relevant curve to realize most of potentially available relative information gain per new data subset (or segment). As Thompson states it, "When the subsets are spaced very closely relative to their length, no information is 'missed' by falling between adjacent subsets, but on the other hand the subsets are highly correlated with each other so that the addition of a subset does not decrease the variance very much." For a highly concentrated window like the 4π prolate spheroidal wave function, about 70% overlap is desirable, according to this criterion. For the second curve down, one would want about 50–60% overlap, and for the third, about 50% would suffice. I have not gone to the effort of generating plots of ΔI_{∞} for other data windows than Thompson's, but one can get a good general impression of where the curves corresponding to other data windows might lie by mentally interpolating between curves, after reference to the bottom portion of Figure 6. For example, the Hanning window is a less highly concentrated data window than the 4π prolate window, but it is more highly concentrated than the "π compound" window, so one might guess that a bit more than 60% overlap would be desirable—this is in accord with Nuttall's tabular data.

References

[1] D. J. Thompson, "Spectrum estimation techniques for characterization and development of WT4 waveguide—I", *Bell System Technical J.*, 56 (1977) 1769–1815; Paper II, pp. 1983–2005, same volume.

[2] A. H. Nuttall, "Spectral estimation by means of overlapped fast Fourier transform processing of windowed data," Technical Report No. 4169, Naval Underwater Systems Center, New London, Conn., October 1971.

[3] A. H. Nuttall, "Estimation of cross-spectra via overlapped fast Fourier transform processing", Technical Report No. 4169-S, Naval Underwater Systems Center, New London, Conn., July 1975.

[4] A. H. Nuttall, "On the variance of the phase estimate of the cross spectrum and coherence", Technical Memorandum No. 771112, Naval Underwater Systems Center, New London, Conn., June 1977.

[5] A. H. Nuttall and G. C. Carter, "A generalized framework for power spectral estimation", *IEEE Trans. Acoust., Speech, Signal Processing*, ASSP-28 (1980) 334–335.

[6] G. M. Jenkins and D. G. Watts, *Spectral Analysis and Its Applications*, Holden–Day, San Francisco, 1968.

[7] L. H. Koopmans, *The Spectral Analysis of Time Series*, Academic Press, New York, 1974.

[8] P. D. Welch, "The use of fast Fourier transform for the estimation of power spectra: a method based on time averaging over short, modified periodograms", *IEEE Trans. Audio and Electroacoust.*, AU-15 (1967) 70–73.

[9] C. Bingham, M. D. Godfrey, and J. W. Tukey, "Modern techniques of power spectrum estimation", *IEEE Trans. Audio and Electroacoust.*, AU-15 (1967) 56–66.

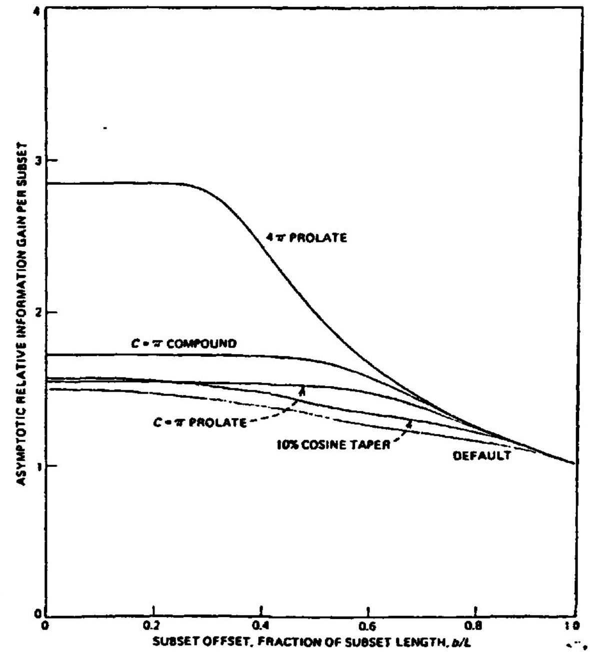

Fig. 8—Asymptotic relative information gain as a function of subset base offset.

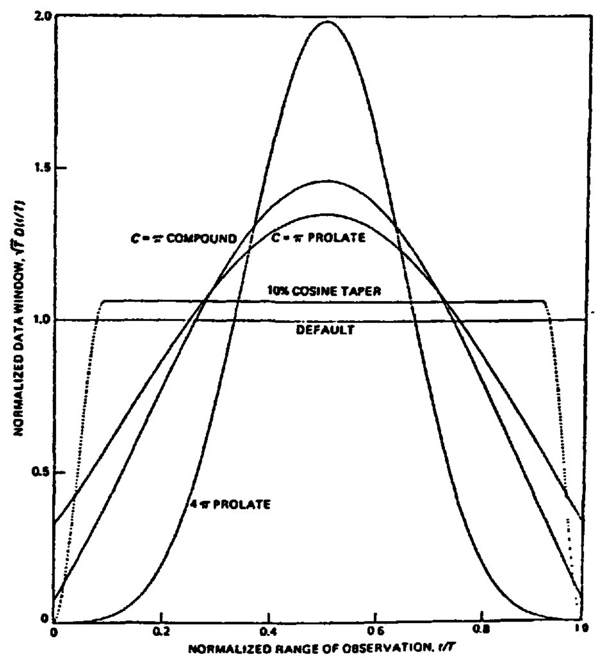

Fig. 1—Comparison of data windows.

Figure 6. In the upper plot are graphs of ΔI_∞ corresponding to the data windows shown in the lower plot. Adapted from [1, Paper I, Figs. 1 and 8].

Appendix A

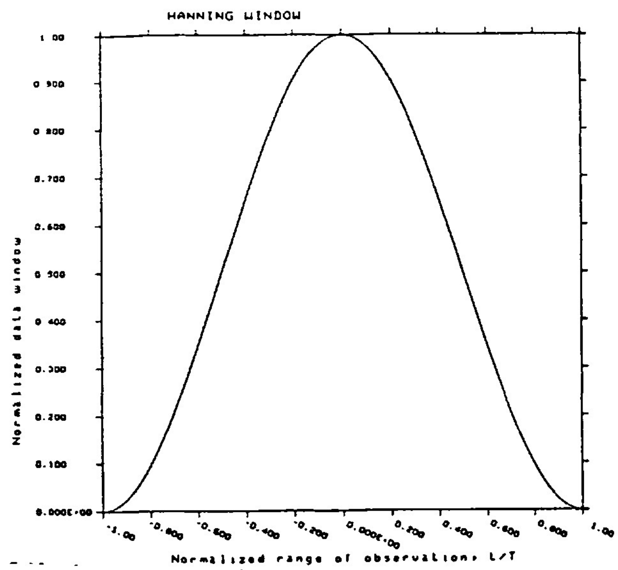
HANNING WINDOW
Normalized data window
Normalized range of observation, t/T

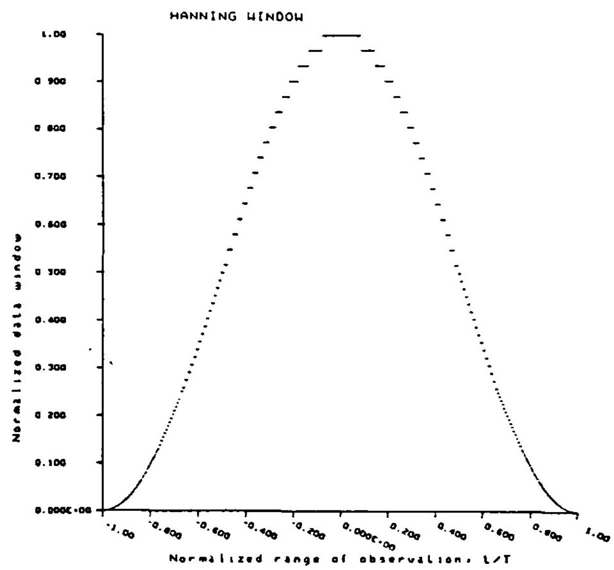
HANNING WINDOW
Normalized data window
Normalized range of observation, t/T

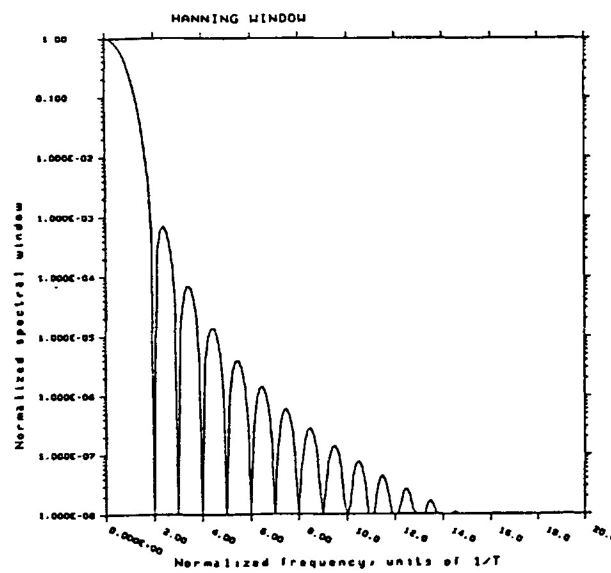
HANNING WINDOW
Normalized spectral window
Normalized frequency, units of 1/T

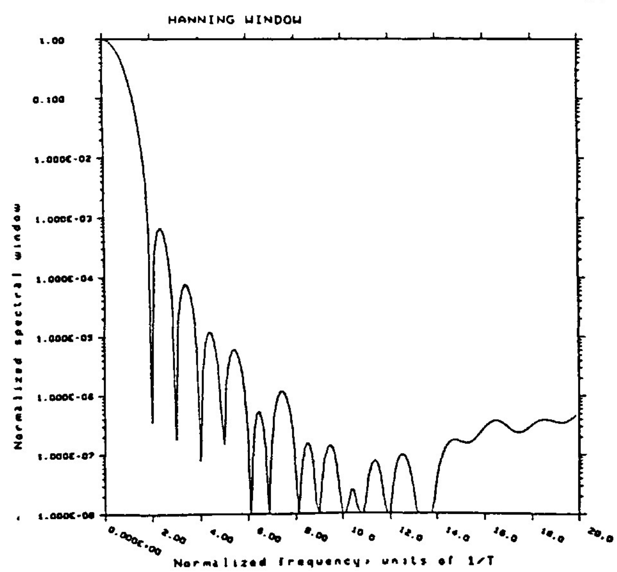
HANNING WINDOW
Normalized spectral window
Normalized frequency, units of 1/T

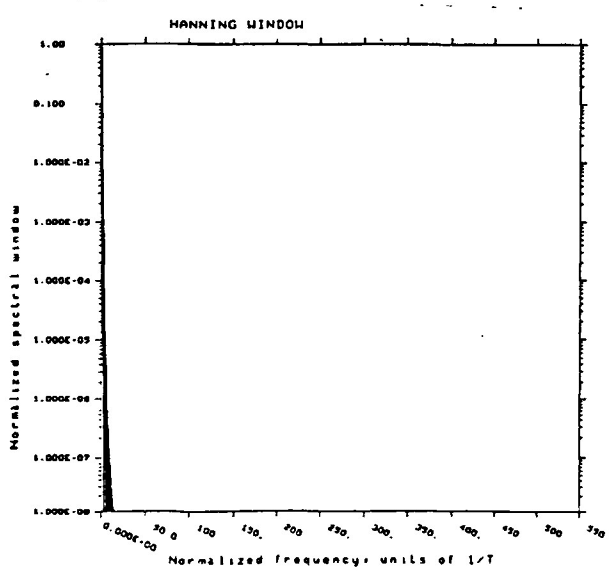
HANNING WINDOW
Normalized spectral window
Normalized frequency, units of 1/T

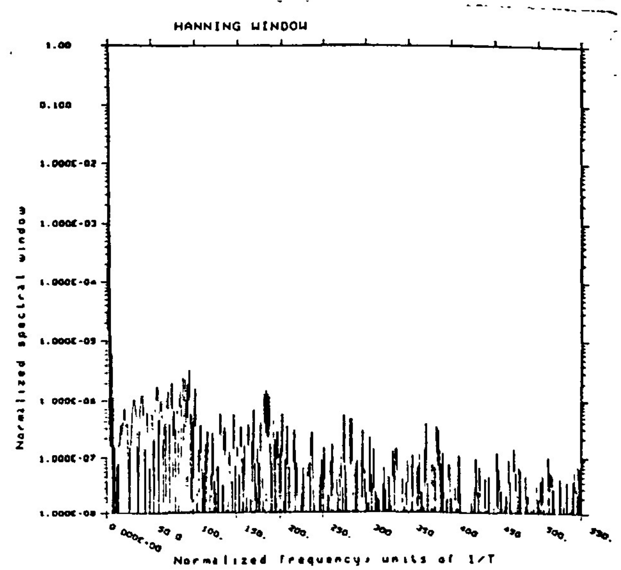
HANNING WINDOW
Normalized spectral window
Normalized frequency, units of 1/T

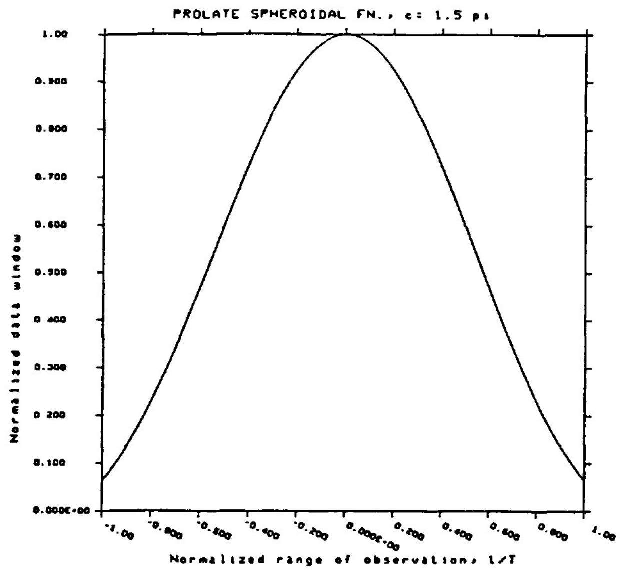
PROLATE SPHEROIDAL FN., c= 1.5 pi
Normalized data window
Normalized range of observation, t/T

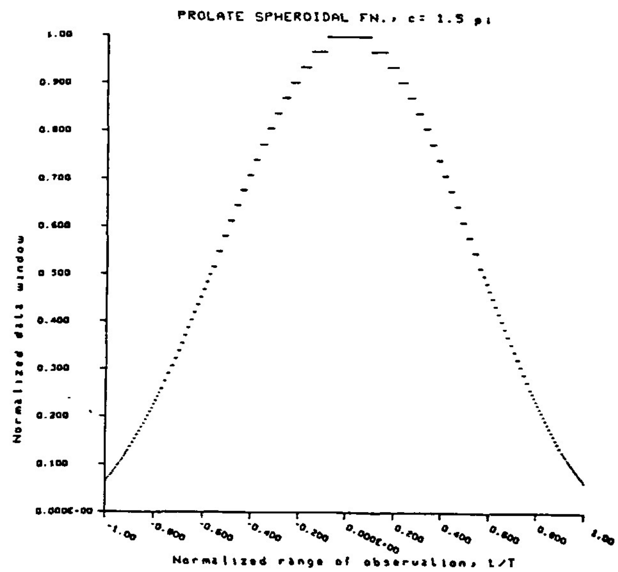
PROLATE SPHEROIDAL FN., c= 1.5 pi
Normalized data window
Normalized range of observation, t/T

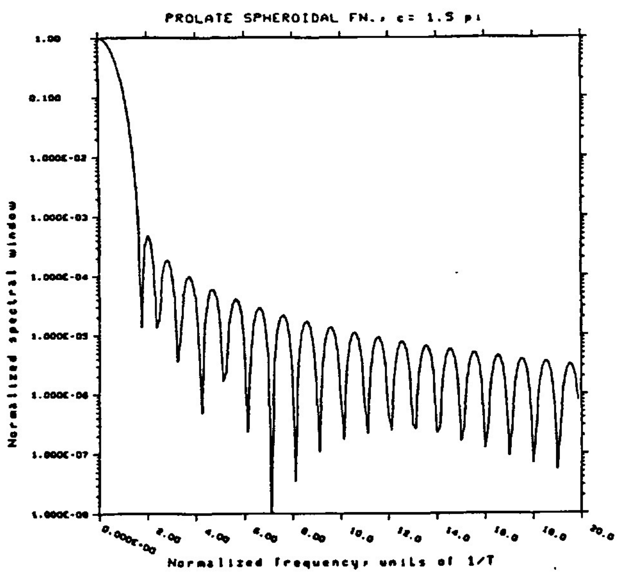
PROLATE SPHEROIDAL FN., c= 1.5 pi
Normalized spectral window
Normalized frequency, units of 1/T

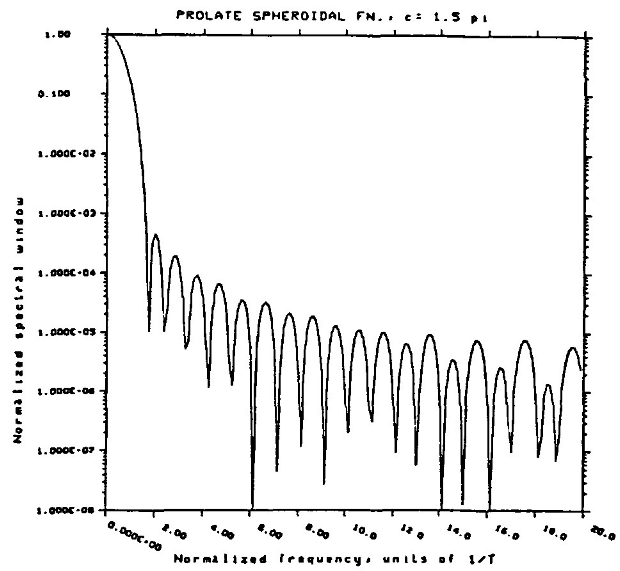
PROLATE SPHEROIDAL FN., c= 1.5 pi
Normalized spectral window
Normalized frequency, units of 1/T

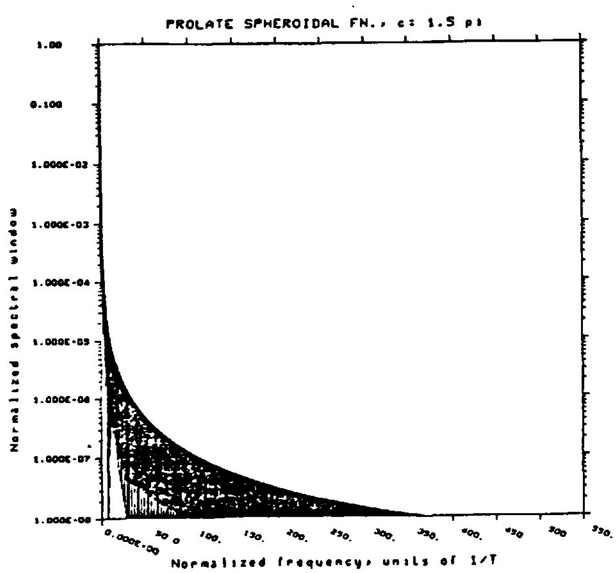
PROLATE SPHEROIDAL FN., c= 1.5 pi
Normalized spectral window
Normalized frequency, units of 1/T

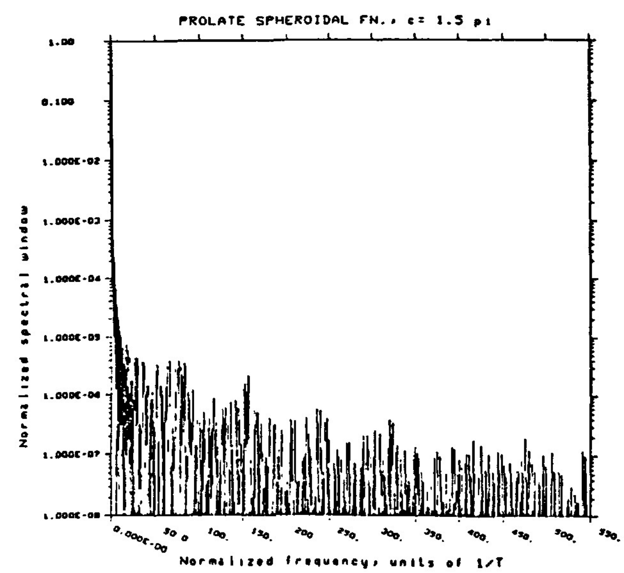
PROLATE SPHEROIDAL FN., c= 1.5 pi
Normalized spectral window
Normalized frequency, units of 1/T

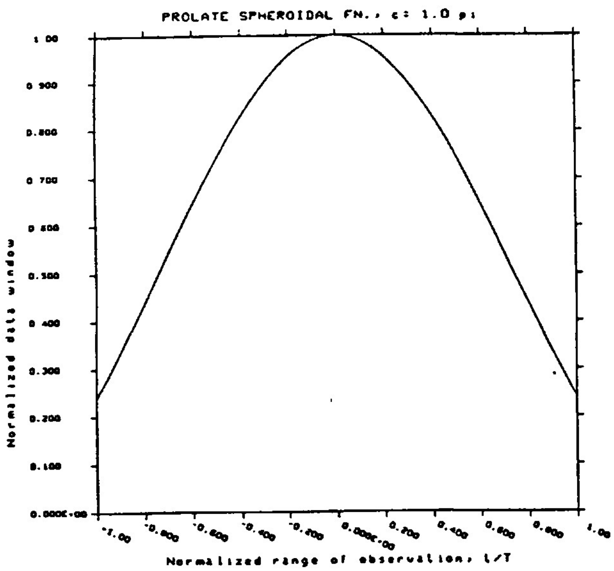
PROLATE SPHEROIDAL FN., c= 1.0 pi
Normalized data window
Normalized range of observation, t/T

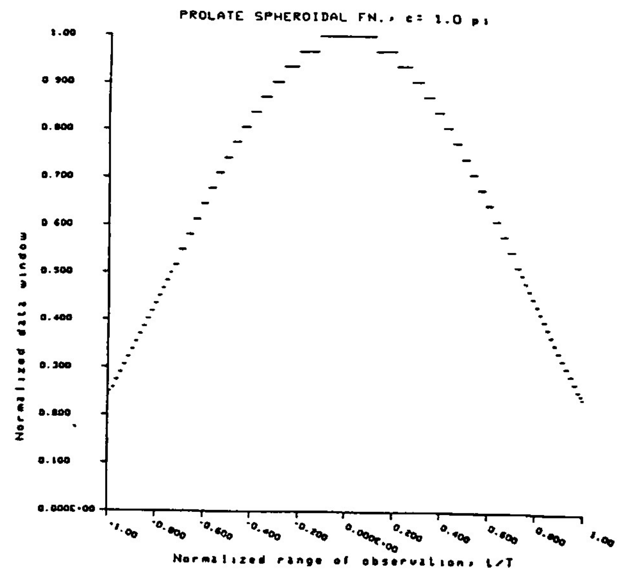
PROLATE SPHEROIDAL FN., c= 1.0 pi
Normalized data window
Normalized range of observation, t/T

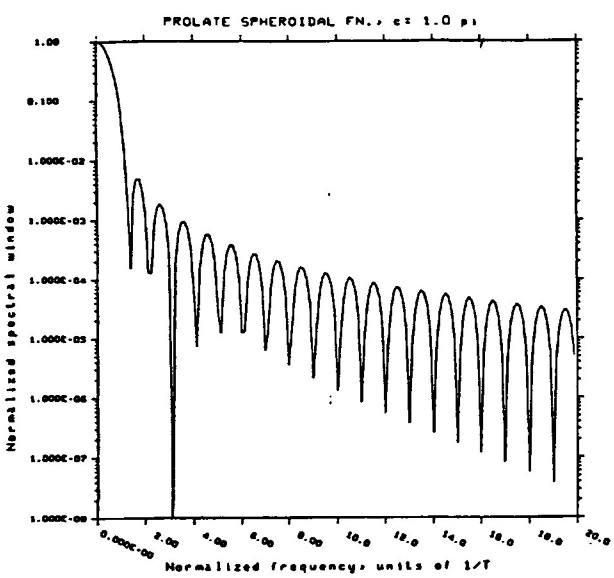
PROLATE SPHEROIDAL FN., c= 1.0 pi
Normalized spectral window
Normalized frequency, units of 1/T

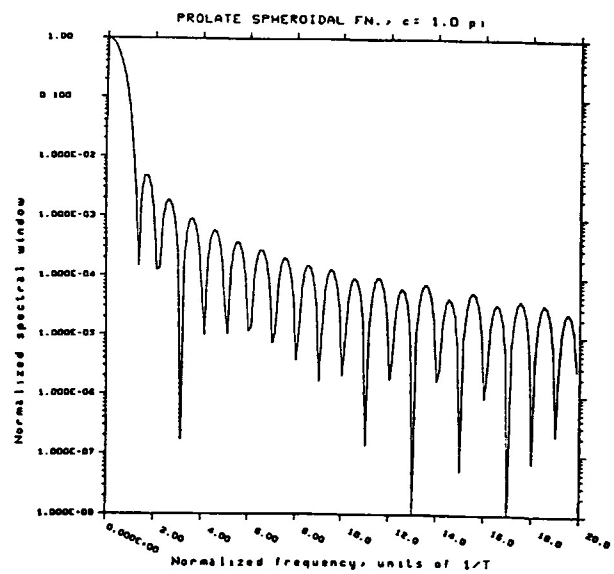
PROLATE SPHEROIDAL FN., c= 1.0 pi
Normalized spectral window
Normalized frequency, units of 1/T

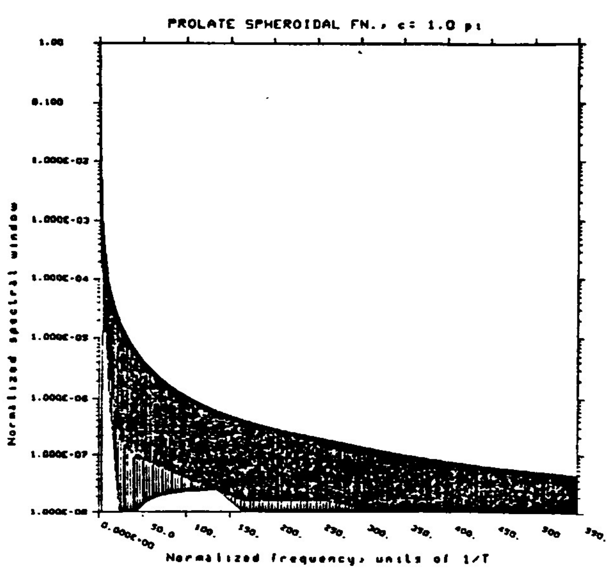
PROLATE SPHEROIDAL FN., c= 1.0 pi
Normalized spectral window
Normalized frequency, units of 1/T

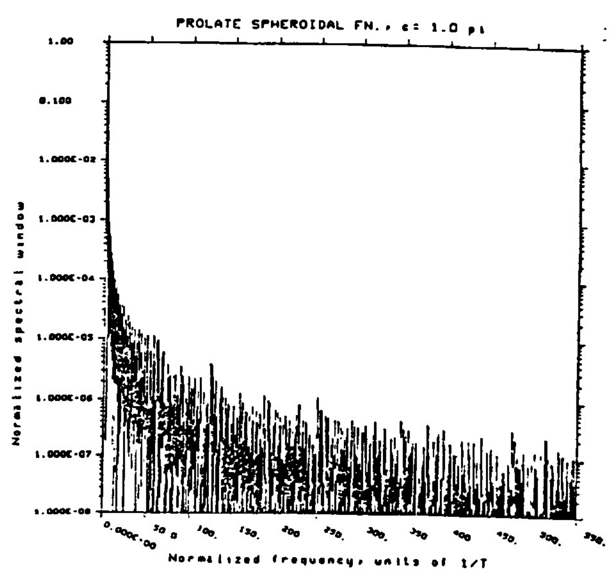
PROLATE SPHEROIDAL FN., c= 1.0 pi
Normalized spectral window
Normalized frequency, units of 1/T

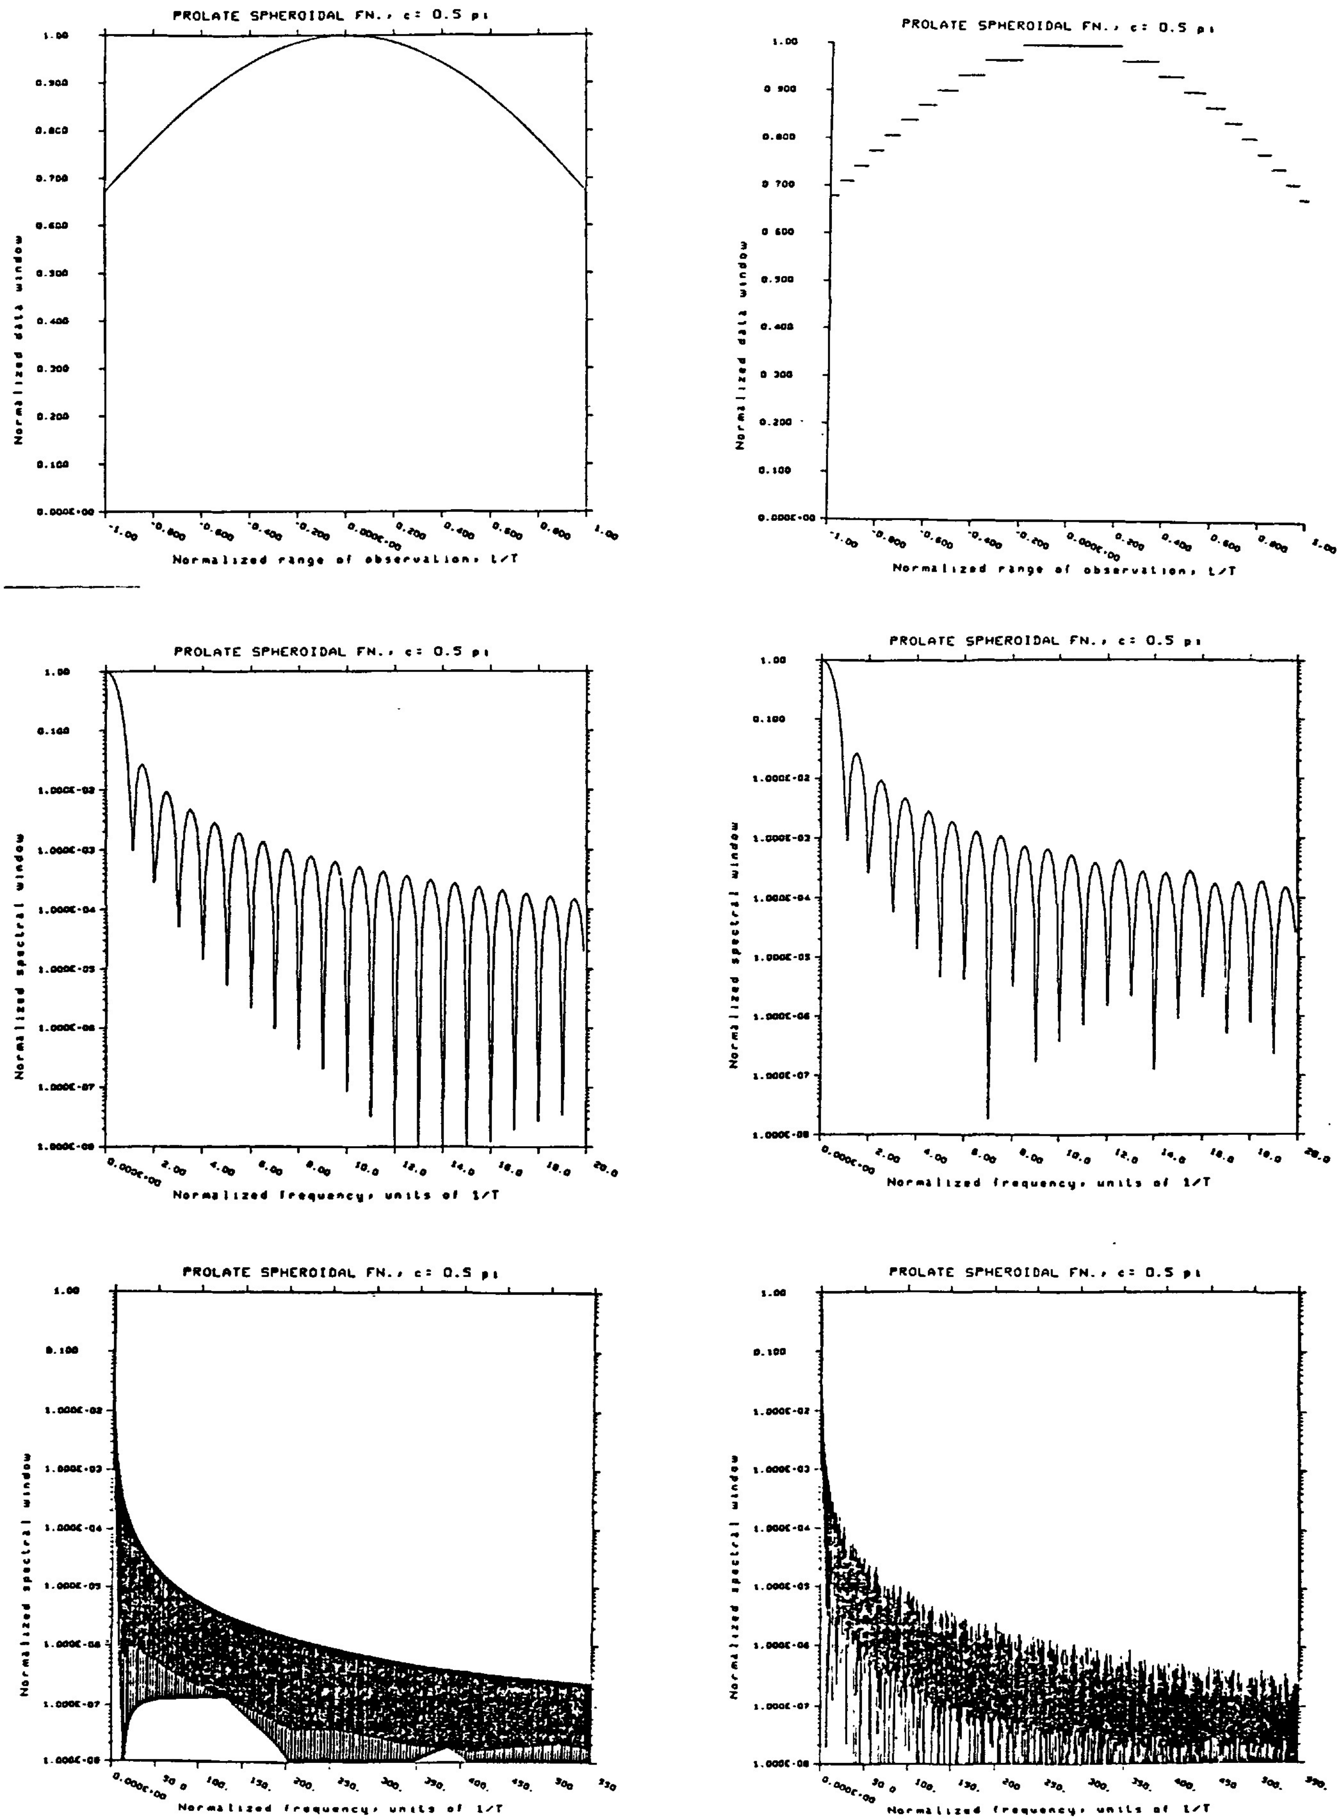
PROLATE SPHEROIDAL FN., c= 0.5 pi
Normalized data window
Normalized range of observation, t/T
Normalized spectral window
Normalized frequency, units of 1/T

Appendix B

NUSC Report No. 4169
13 October 1971

Spectral Estimation by Means of Overlapped Fast Fourier Transform Processing of Windowed Data

Albert H. Nuttall

ABSTRACT

An investigation of power-density autospectrum estimation by means of overlapped Fast Fourier Transform (FFT) processing of windowed data is conducted for four candidate spectral windows with good side-lobe behavior. A comparison of the four spectral windows is made on the basis of equal half-power resolution bandwidths. The criteria for comparison are: (1) statistical stability of the spectral estimates, (2) leakage (side lobes) of the spectral windows, (3) number of FFTs (number of overlapped pieces) required, and (4) size of each FFT required. The dependence of these criteria on the amount of overlap is investigated quantitatively.

Some striking invariances are discovered. Specifically, it is shown that the ultimate variance-reduction capabilities of the four windows, as measured by the equivalent number of degrees of freedom (EDF), are virtually identical under the constraint of equal half-power bandwidths. Furthermore, when the proper overlap is used for each window, the stability of this method of spectral estimation is identical to that of the "indirect" correlation approach. Also, the number of FFTs required to realize 99 percent (or less) of the maximum EDF is virtually independent of the particular window employed. The required fractional overlap of the four data windows for 99 percent (or less) of the maximum EDF is virtually independent of the product of the available time and the resolution bandwidth, although it does depend on the particular window. Tables of required overlap are presented. The only tradeoff among the four windows is that those with better side lobes require larger-size FFTs. All of these results are derived for a Gaussian random process, under the assumption that the resolution bandwidth of the spectral window is smaller than the finest detail in the true spectrum.

Rules of thumb for the maximum EDF and the number of FFTs required to realize 99 percent of the maximum EDF are given. The possibility of weighting individual spectral estimates unequally in order to optimize the EDF is investigated; the gain is found to be negligible for cases of practical interest.

NUSC Technical Report 4169-S
11 July 1975

Estimation of Cross-Spectra Via Overlapped Fast Fourier Transform Processing

Albert H. Nuttall

ABSTRACT

The optimum overlap to be used for estimation of cross-spectra via FFT processing of windowed data is shown to be identical to that for estimation of auto-spectra. In addition, a useful geometric interpretation of the random errors in cross-spectral estimation, and their covariances, is furnished.

B-2

Technical Memorandum 771112
10 June 1977

On the Variance of The Phase Estimate of The Cross Spectrum And Coherence

A. H. Nuttall

ABSTRACT

The variance of the phase estimate of the cross spectrum and coherence is numerically evaluated for values of the true magnitude-squared coherence, S, equal to 0(.1).9 and .99, and for the number of independent averages, n, equal to 1(1)500. It is found that the approximation (1-S)/(SK), where K = 2n for independent averages, is a good one for all S and for K > 10, although the approximation is generally optimistic. A useful recursion formula for the probability density function of the phase estimate is also derived. The danger of employing a Gaussian approximation is demonstrated dramatically in a numerical example. An extension of the equivalent degrees of freedom to complex averages is made and suggested for use in cross-spectral estimation.

NUSC Report 5529
3 December 1976

Probability Distribution of Spectral Estimates Obtained Via Overlapped FFT Processing of Windowed Data

Albert H. Nuttall

ABSTRACT

The characteristic function of spectral estimates obtained via overlapped FFT processing of windowed data is presented for a random process containing a signal tone and Gaussian noise. For the special case of noise-alone, the probability distribution of the estimate is plotted and compared with an approximation utilizing only the first two moments and found to be in excellent agreement in probability over the range (.0001, .9999) for several data windows, overlaps, and time-bandwidth products. This result means that knowledge of the equivalent degrees of freedom of the spectral estimate is adequate for a complete probabilistic description, even when the overlap results in significant statistical dependence of the component FFT outputs.

B-4