

Spectrum and spectral density estimation by the Discrete Fourier transform (DFT), including a comprehensive list of window functions and some new flat-top windows.

G. Heinzel*, A. Rüdiger and R. Schilling,
Max-Planck-Institut für Gravitationsphysik
(Albert-Einstein-Institut)
Teilinstitut Hannover

February 15, 2002

Abstract

This report tries to give a practical overview about the estimation of power spectra/power spectral densities using the DFT/FFT. One point that is emphasized is the relationship between estimates of power spectra and power spectral densities which is given by the effective noise bandwidth (ENBW). Included is a detailed list of common and useful window functions, among them the often neglected flat-top windows. Special highlights are a procedure to test new programs, a table of comprehensive graphs for each window and the introduction of a whole family of **new flat-top windows** that feature sidelobe suppression levels of up to -248 dB, as compared with -90 dB of the best flat-top windows available until now.

Contents

1	Abbreviations and symbols	4
2	Preface	5
3	Introduction	5
4	Scaling bits to Volts	7
5	Discrete Fourier Transformation (DFT)	7
6	Some common units	9
6.1	peak, peak-peak and rms	9
6.2	decibel	10
7	Sampling time, frequency bins etc.	10

*e-mail: ghh@mpq.mpg.de

8 Window functions	11
8.1 The Hanning window as an example	11
8.2 The window function in the frequency domain	14
9 Scaling the results	15
10 Averaging and overlap	17
11 Preprocessing of the data: DC and trend removal	20
12 Putting it all together	21
13 Testing the algorithm	21
14 MATLAB pwelch function	23
Appendices	25
A Maximal dynamic range of real signals	25
A.1 Electronic noise	25
A.2 Digitizing noise	25
B Computing the response of a window function	27
C List of window functions	28
C.1 Rectangular window	30
C.2 Bartlett window	30
C.3 Welch window	31
C.4 Hanning window	31
C.5 Hamming window	32
C.6 Blackman-Harris window	32
C.7 Nuttall windows	33
C.7.1 Nuttall3	34
C.7.2 Nuttall3a	35
C.7.3 Nuttall3b	35
C.7.4 Nuttall4	35
C.7.5 Nuttall4a	35
C.7.6 Nuttall4b	36
C.7.7 Nuttall4c	36
C.8 Kaiser window	36

D	Flat-Top windows	38
D.1	Published flat-top windows	39
D.1.1	Fast decaying 3-term flat top window	40
D.1.2	Fast decaying 4-term flat top window	40
D.1.3	Fast decaying 5-term flat top window	40
D.1.4	Minimum sidelobe 3-term flat top window	41
D.1.5	Minimum sidelobe 4-term flat top window	41
D.1.6	Minimum sidelobe 5-term flat top window	41
D.2	Flat-top windows by instrument makers	42
D.2.1	National Instruments flat-top window	42
D.2.2	Old HP flat-top window	42
D.2.3	Undisclosed HP flat-top window	43
D.2.4	Stanford Research flat-top window	43
D.3	New flat-top windows	45
D.3.1	HFT70	45
D.3.2	HFT95	46
D.3.3	HFT90D	46
D.3.4	HFT116D	47
D.3.5	HFT144D	47
D.3.6	HFT169D	47
D.3.7	HFT196D	48
D.3.8	HFT223D	48
D.3.9	HFT248D	49
E	Bibliography and Acknowledgements	49
F	Graphs of window functions	50

1 Abbreviations and symbols

The following list contains those symbols and abbreviations that appear in more than one place in the text.

$a(f)$	Fourier transform (‘transfer function’) of a window in the frequency domain, see Section 8.2 and Appendix B
c_k	coefficients of cosine terms in definition of window function
$e_{\max}$	maximal amplitude error in one bin ($-0.5 \leq f \leq 0.5$)
f_{res}	frequency resolution (width of one frequency bin), see Section 7
f_s	sampling frequency, see Section 7
n	number of bits of an ADC
N	Length of the DFT/FFT
S_1	sum of window values used for normalization, see Equation (19)
S_2	sum of squared window values used for normalization, see Equation (20)
U_{LSB}	smallest voltage step of an ADC, corresponding to one ‘least significant bit’
w_j	j -th member of the window function vector
x_j	digitized time series; input to DFT in time domain
y_m	raw output of DFT in frequency domain
ADC	Analog-to-Digital Converter
AF	Amplitude Flatness, see Section 10
BW	BandWidth
dB	deciBel, see Section 6
DC	‘Direct Current’, constant component of a signal
DFT	Discrete Fourier Transform
ENBW	Effective Noise BandWidth, see Equation (22)
FFT	Fast Fourier Transform
FFTW	A software package that implements the FFT
GPL	Gnu Public License
LS	Linear (amplitude) Spectrum
LSB	Least Significant Bit
LSD	Linear Spectral Density
MATLAB	– Commercial software package –
NENBW	Normalized Equivalent Noise BandWidth, see Equation (21)
OC	Overlap Correlation, see Section 10
PF	Power Flatness, see Section 10
PS	Power Spectrum
PSD	Power Spectral Density
PSLL	Peak SideLobe Level
rms	root mean square, see Section 6
ROV	Recommended Overlap, see Section 10
SLDR	SideLobe Drop Rate
WOSA	Welch’s Overlapped Segmented Average

2 Preface

The definition and usage of the Fourier transform as it is widely used, e.g., in theoretical physics considerably differs from the practical application of the Discrete Fourier Transform (DFT) in data analysis.

The subject of this report is the estimation of spectra and spectral densities using the DFT. While many methods for spectrum estimation are discussed in the statistical literature, we present here only the one method that has found the widest application in engineering and experimental physics: the ‘overlapped segmented averaging of modified periodograms’. A **periodogram** here means the discrete Fourier transform (DFT) of one **segment** of the time series, while **modified** refers to the application of a time-domain window function and **averaging** is used to reduce the variance of the spectral estimates. All these points will be discussed in the following sections. This method is attributed to Welch[1] and is also known under various acronyms such as WOSA (for ‘Welch’s overlapped segmented average’) etc. While rather straightforward in theory, the practical implementation involves a number of nontrivial details that are often neglected. This report tries to clarify some of these problems from a practical point of view. It was written for the needs of the data analysis of the laser-interferometric gravitational wave detector GEO 600, but is expected to be useful for many other applications as well.

It only treats the computation of a spectrum or spectral density, starting from a digitized time series, typically measured in Volts at the input of the A/D-converter. The conversion of the resulting spectrum into other units (such as meters) and related problems such as un-doing the effect of whitening filters are not treated here.

This report is organized as follows: In Sections 3 to 11, individual problems such as scaling, windowing, averaging etc. are discussed. An overview of how to combine all these details into a working algorithm is given in Section 12. Two somewhat more specialized topics are treated in Appendices A and B. The Appendices C and D describe in detail many common and useful window functions, including several new high-performance functions recently developed by one of the authors (G.H.). An overview of all window functions in tabular form is given on page 29. The final part of this report, Appendix F starting on page 50, consists of graphs showing the most important characteristics of each window function.

3 Introduction

As an example, we consider a signal $x(t)$ that contains a sinusoid at 1234 Hz with an amplitude of $2 V_{\text{rms}}$. Furthermore it contains white noise with a density of $10 \text{ mV}_{\text{rms}}/\sqrt{\text{Hz}}$, band-limited to the range between 50 Hz and 2000 Hz. The spectrum of that signal is sketched in Figure 1, with the unit of the y -axis intentionally not yet well defined.

Such a signal is quite realistic. It can easily be produced by electronic circuits and it is not untypical for the signals that we might expect in the data analysis of laser-interferometric gravitational wave detectors such as GEO 600, just simpler.

Let us further assume that this signal is sampled by an A/D-converter with a constant sampling frequency f_s of, e.g., 10 kHz. We thus have a time-series of equidistant samples from the A/D-converter as input to our data processing task.

The desired output is a graph resembling Figure 1 which allows us to determine both the amplitude of the peak ($2 V_{\text{rms}}$) and the level of the noise floor ($10 \text{ mV}_{\text{rms}}/\sqrt{\text{Hz}}$). As can already be seen from the different units (V vs. $\text{V}/\sqrt{\text{Hz}}$) we will need different scalings of the y -axis for




Figure 1 Sketch of the example spectrum to be discussed in the text. Both axes are plotted logarithmically. The unit of the y -axis is intentionally not yet well defined.

the two cases. It will turn out that the ratio between the two different scaling factors (the square root of the effective noise bandwidth) depends on internals of the algorithm that is used and must thus be transparent in the final result. Hence the desired output are in fact **two** graphs, one scaled in V and another one in $V/\sqrt{\text{Hz}}$. We will call the former one a (linear) **spectrum** and the latter a (linear) **spectral density**.

The power spectral density describes how the power of a time series is distributed with frequency. Mathematically, it is defined as the Fourier transform of the autocorrelation sequence of the time series.

The *linear* spectral density is simply the square root of the *power* spectral density, and similarly for the spectrum. In GEO 600 the linear spectral density, which has a unit such as $V/\sqrt{\text{Hz}}$, is used very often. It is usually indicated by placing a tilde ($\sim$) over the symbol for the quantity in question, i.e. the linear spectral density corresponding to a voltage $U_{\text{something}}(t)$ is represented by the symbol $\tilde{U}_{\text{something}}(f)$. One practically important relation for a linear spectral density $\tilde{U}(f)$ is its relation to the rms fluctuation of the quantity U , assumed to be band-limited to the frequency range $f_1 \leq f \leq f_2$:

$$U_{\text{rms}} = \sqrt{\int_{f_1}^{f_2} [\tilde{U}(f)]^2 df}, \quad (1)$$

Table 1 summarizes the 4 different ways to express the result of the DFT, assuming the input time series is measured in Volts.

The relationship between continuous and discrete Fourier transforms, the difference between one-sided and two-sided spectra etc. are beyond the scope of this report, but are discussed in many books (see also Reference [Schutz]). However, we don't need these concepts for the rest of this report, where we are concerned with the practical computation only.

Experimentalists often prefer the linear versions because they can be directly related to the parameters of the experiment. For example, if a signal generator that produces a sine wave with an amplitude of $2V_{\text{rms}}$ at 1234 Hz is connected to the A/D converter, we expect in the resulting linear spectrum a peak at the right frequency, the height of which directly corresponds to the amplitude that we had set on the signal generator. This requirement alone already establishes the only correct scaling for linear spectra from an experimentalist's point of view.

The relationship between a spectrum and a spectral density is given by the effective noise bandwidth and can easily be determined at the time when the DFT is computed. However, it

Abbrev.	Name	Relation	Unit
PSD	power spectral density		V ² /Hz
PS	power spectrum	PS = PSD × ENBW	V ²
LSD	linear spectral density amplitude spectral density	LSD = √PSD	V/√Hz
LS	linear spectrum amplitude spectrum	LS = √PS = LSD × √ENBW	V

Table 1 Naming convention for DFT outputs. ENBW is the equivalent noise bandwidth, defined in Equation (22)

cannot normally be reproduced from either of the two graphs without additional information.

4 Scaling bits to Volts

Since the data that we want to analyze are measured in Volts at the input of the A/D-converter (ADC) and have already been converted into binary bits when we get them, the first scaling factor that usually must be applied to the raw data is the conversion from bits back to Volts. This is usually straightforward: if the data arrive as binary signed integers, they are converted to floating-point numbers by multiplication with the factor

$$U_{\text{LSB}} = \frac{U_{\text{max}} - U_{\text{min}}}{2^n}, \quad (2)$$

where U_{LSB} is the voltage corresponding to a least significant bit and n the number of bits of the A/D-converter. Gain and offset of any preamplifiers used directly before the A/D-converter can also be applied at this stage.

U_{max} and U_{min} are *not* the outer limits of the input voltage range, but rather two common points of the highest and lowest voltage bin, i.e. either their centers, or their maxima or their minima. This detail is unimportant for big n .

5 Discrete Fourier Transformation (DFT)

The discrete Fourier transform in general takes a vector of N complex numbers x_k , $k=0 \dots N-1$, and transforms it into a vector of N complex numbers y_m , $m=0 \dots N-1$.

There are three different common definitions for the DFT, differing only in their normalization

factors:

$$y_m^{(1)} = \sum_{k=0}^{N-1} x_k \exp\left(-2\pi i \frac{mk}{N}\right), \quad m = 0 \dots N-1, \quad (3)$$

$$y_m^{(2)} = \frac{1}{\sqrt{N}} y_m^{(1)}, \quad (4)$$

$$y_m^{(3)} = \frac{1}{N} y_m^{(1)}. \quad (5)$$

This transformation is called the ‘forward’ DFT, with its inverse, the ‘backward’ DFT, differing in the sign of the exponent ($+2\pi i \frac{mk}{N}$ instead of $-2\pi i \frac{mk}{N}$).

The practical implementation of the DFT on a computer nearly always uses the Fast Fourier Transform (FFT) algorithm, and the terms DFT and FFT are used interchangeably in this report. Most FFT subroutine packages (e.g. [FFTW]), compute $y^{(1)}$.

The `Fourier` function of MATHEMATICA computes $y^{(2)}$, which incidentally is the only ‘symmetrical’ way to define forward and backward DFT such that successively applying both the forward and then the backward DFT reproduces the original data.

For the simplest computation of the spectrum, however, we need $y^{(3)}$. That can be directly seen if we consider the sinusoidal signal again. If we increase N , $y^{(1)}$ will pick up more and more of the signal and the result in the adequate frequency bin will increase proportionally to N , while our desired result, the amplitude, obviously must not depend on N . If, as is usual, window functions are being used, the window-related normalization will take care of this factor (see Section 9). Nevertheless, when an FFT algorithm is used for the first time, care should be taken to determine which of (3), (4) or (5) it actually computes.

While the x_k are complex numbers in the general definition of the DFT, the time series of the digitized input signal is always real. As a consequence, the output array y_m obeys the following relationship:

$$y_{N-m} = y_m^*, \quad (6)$$

where $*$ denotes complex conjugation. In particular, if N is even, y_0 and $y_{N/2}$ are real. From now on we will assume that N is even. The upper half of the array, $y_{N/2+1} \dots y_{N-1}$ is thus redundant and never computed by efficient FFT packages. The results are typically returned in some ‘packed’ format. For the case of the FFTW package (see Reference [FFTW]), the result is a real vector z_i , $i = 0 \dots N-1$ with

$$z_i = \Re\{y_i\}, \quad 0 \leq i \leq N/2, \quad (7)$$

$$z_{N-i} = \Im\{y_i\}, \quad 0 < i < N/2. \quad (8)$$

If only a complex-to-complex FFT subroutine is available, Section 12.3 of Reference [Press] describes two ways of using it to compute real-to-complex transforms with moderate efficiency.

Because of the above relationship it is convenient to have N even, as we do assume throughout this report. Some subroutine packages even demand that N be a power of 2. That is, however, an unnecessarily stringent limitation in many situations. The FFTW package, for example, will compute the DFT for any positive integer N , and with high efficiency for integers of the form

$$N = 2^a 3^b 5^c 7^d 11^e 13^f, \quad (9)$$

where $e + f$ is either 0 or 1 and the other exponents are arbitrary. This gives more flexibility in the choice of the frequency resolution f_{res} , see Sections 7 and 12 below. Furthermore, it can also be configured to efficiently compute the FFT for other integers, if necessary.

6 Some common units

6.1 peak, peak-peak and rms

The next simple scaling factor concerns the units of the amplitude (Volts in this example). This matter equally concerns both the spectra and the spectral densities. Three different units are in common usage (see Figure 2):




Figure 2 *Three ways to measure an amplitude.*

V_{pk} ‘Volts peak’ is the most obvious measurement of an amplitude, it is the height of the peak measured from the average value (i.e. from zero if there is no DC component).

V_{pk-pk} ‘Volts peak-to-peak’. This measurement is exactly two times bigger than ‘Volts peak’ and popular with experimentalists because it can easily be read directly from an oscilloscope (regardless of any DC offsets).

V_{rms} ‘Volts rms’ or ‘Volts root-mean-square’. This measurement is defined as the DC voltage that has the same *power* as the signal in question. For sinusoidal signals (all signals at the output end of a DFT are sinusoidal) the relationship is¹

$$V_{rms} = \frac{V_{pk}}{\sqrt{2}}. \quad (10)$$

In the following we will use V_{rms} as primary unit for all DFT results. If necessary they can be converted to the other units as described above.

While the output of the DFT at a certain frequency will be a complex number, the amplitude is defined as the absolute value of that complex number, i.e. amplitudes are non-negative real numbers.

¹While in general V_{rms} is defined via an integral and the result depends on the precise limits of that integral, in this context (as an amplitude measurement for sinusoidal signals) it is always assumed that an integer number of cycles is averaged.

6.2 decibel

To complicate matters even further, all of the above three measurements can be expressed in ‘decibel’. The general definition of decibel (dB) is:

$$\text{ratio[dB]} = 10 \log_{10} \left(\frac{\text{signal power}}{\text{reference power}} \right) = 20 \log_{10} \left(\frac{\text{signal amplitude}}{\text{reference amplitude}} \right). \quad (11)$$

Any value given in ‘dB’ hence expresses a ratio and not an absolute amplitude. The reference amplitude is often expressed after the symbol ‘dB’, e.g. ‘dBV’ or ‘dBμV’ which refer to a reference amplitude of 1 V or 1 μV, respectively.

In high-frequency electronics the most common measurement of powers and amplitudes is called ‘dBm’. The reference here is 1 mW of power delivered into the nominal load impedance, usually 50 Ω.

The conversion of an amplitude x given in dBm to the amplitude voltage U expressed in V_{rms} is given by:

$$U_{\text{rms}} = \sqrt{P \times 50 \Omega}, \quad (12)$$

where the power P is computed by

$$1 \text{ W} \times 10^{\{-3 + 0.1 \times \frac{x}{1 \text{ dBm}}\}} \quad (13)$$

As an example, the amplitude of the 1234 Hz component of our test signal could be expressed as

$$2 V_{\text{rms}} = 2.828 V_{\text{pk}} = 5.657 V_{\text{pk-pk}} = 6.02 \text{ dBV}_{\text{rms}} = 126.02 \text{ dB}\mu V_{\text{rms}} = 19.03 \text{ dBm}. \quad (14)$$

7 Sampling time, frequency bins etc.

While the DFT transforms N numbers into N new numbers and does not care about sampling times, frequency bins etc., these notions are essential for a correct interpretation of its results.

Assume the sampling frequency f_s to be fixed. From the Nyquist theorem it follows that the maximal useful frequency is $f_{\text{Ny}} = f_s/2$. Then the only degree of freedom that we have left in the data processing is the length N of the DFT.

The width of a frequency bin (also called frequency resolution) is given by

$$f_{\text{res}} = \frac{f_s}{N}. \quad (15)$$

Note that the ability to separate two nearby peaks in the spectrum is also sometimes called frequency resolution, but not in this report. It depends strongly on the window function that was chosen (in particular on the bandwidth of the window, see Section 8 below). The minimal separation of two signals that can just be separated also depends on the ratio between the two amplitudes and is typically a few times larger than our f_{res} , see for example the last Section of Reference [Harris78].

As we have seen above, the DFT will produce only $N/2 + 1$ distinct complex numbers. These correspond to the frequencies

$$f_m = m f_{\text{res}}, \quad m = 0 \dots N/2. \quad (16)$$

The first element, y_0 , corresponds to the DC average of the signal and accordingly has no imaginary part. The last element, $y_{N/2}$, corresponds to the Nyquist frequency $f_{N/2} = (N/2) f_{\text{res}} = f_s/2 = f_{\text{Ny}}$ and also has no imaginary part² because of the complex-conjugate symmetry of the output array. The first and last frequency bins are usually not very useful anyway (see Section 9).

Hence in a practical approximation we can say that we obtain $N/2$ frequency bins of the width f_s/N in the output. The apparent loss of information is explained by the fact that the output consists of complex numbers while the input was real.

The phase given by the complex numbers in the result is random for those frequency bins that are dominated by stochastic noise; it may be meaningful for frequency bins that are dominated by a sinusoidal signal. In that case, it is typically used by processing a reference sinusoidal signal of the same frequency with the same algorithm and subtract the two resulting phases to obtain the relative phase between the two sinusoidal signals while canceling any phase offset due to the algorithm. In most applications, however, the phases resulting from the DFT are ignored by immediately taking the absolute values, such as in Equations (23) or (24).

8 Window functions

If we simply take a stretch of length N out of a time series containing a sinusoidal signal and perform a DFT, we will most likely find that the sinusoidal signal which we might naively expect to result in a sharp peak in only one frequency bin, will instead show up as something ugly like shown in Figure 3.

The reason is that the DFT implicitly assumes that the signal is periodic, i.e. that the time series of length N repeats itself infinitely in a cyclic manner. If the frequency of the sinusoidal input signal is not an exact multiple of the frequency resolution f_{res} , i.e. does not fall in the exact center of a frequency bin, this assumption is not true, and the DFT will ‘see’ a discontinuity between the last sample and the first sample due to the cyclic continuation. That discontinuity spreads power all across the spectrum.

The remedy is to multiply the time series with a ‘window function’ in the time domain before applying the DFT. This window function starts near or at zero, then increases to a maximum at the center of the time series and decreases again. Thus the discontinuity is removed. Many window functions have been defined and given names. They usually involve some compromise between the width of the resulting peak in the frequency domain, the amplitude accuracy and the rate of decrease of the spectral leakage into other frequency bins.

8.1 The Hanning window as an example

This section discusses the general properties of window functions, using the simple but useful Hanning window as an example. Other window functions, including flat-top windows, are discussed in the Appendices C and D.

A window function to be used with a DFT of length N is defined by a vector of real numbers $\{w_j\}$, $j = 0 \dots N - 1$. It is used by multiplying the time series x_j with the window before performing the DFT, i.e. using $x'_j = x_j \cdot w_j$ as input to the DFT.

²If N is even as we assume here.




Figure 3 *The spectral response of a rectangular window.*

The Hanning window, our example in this Section, is defined as follows:

$$w_j = \frac{1}{2} \left[1 - \cos \left(\frac{2\pi \cdot j}{N} \right) \right]; \quad j = 0 \dots N - 1, \quad (17)$$

and is shown in Figure 4.

All windows studied here have the following symmetry:

$$w_j = w_{N-j}. \quad (18)$$

This implies for even lengths N that w_0 and $w_{N/2}$ appear only once, while all other coefficients appear twice. Hence only the $N/2 + 1$ coefficients $w_0, \dots, w_{N/2}$ need to be computed and stored. Note also that this symmetry implies that the N -element vector $\{w_j\}, j = 0 \dots N - 1$ is not as symmetric as Figure 4 might suggest, i.e. in general we have $w_0 \neq w_{N-1}$. The symmetry relation given in Equation (18) is illustrated in Figure 5 for $N = 8$.

We define the following two sums for normalization purposes:

$$S_1 = \sum_{j=0}^{N-1} w_j, \quad (19)$$

$$S_2 = \sum_{j=0}^{N-1} w_j^2. \quad (20)$$

Because we will use S_1 and S_2 in the normalization of our final results, we can multiply the window values w_j with any convenient constant factor.




Figure 4 *The Hanning window in the time domain.*




Figure 5 *The symmetry relation (18) illustrated for a Hanning window with $N = 8$. The window is shown applied sequentially three times.*

8.2 The window function in the frequency domain

The transfer function $a(f)$ of a window expresses the response of the window to a sinusoidal signal at an offset of f frequency bins. It can be computed from the window values w_j according to the formulae given in Appendix B.

The transfer function $a(f)$ of the Hanning window is shown in Figure 6. Apart from the central peak (which we want) there are many more peaks at regular intervals. They are called ‘sidelobes’ and are undesirable. One aim in designing window functions is to reduce the height of these sidelobes. Note the different y -axes in Figures 3 and 6; at a frequency offset of 20 bins the sidelobes of the rectangular window are at -36 dB, while those of the Hanning window are at -88 dB.




Figure 6 *The Hanning window in the frequency domain.*

The 3 dB bandwidth $W_{3\text{dB}}$ that is listed in Appendix C for each window is computed from Equation (42) by numerically solving the equation $a(f) = -3$ dB for f . For the Hanning window it is $W_{3\text{dB}} = 1.44$ bins. Of course one wants this bandwidth not to be unnecessarily wide. It turns out, however, that reducing the level of the sidelobes increases the bandwidth, such that a compromise must be found.

The normalized equivalent noise bandwidth NENBW of the window, expressed in frequency bins, is given by

$$\text{NENBW} = N \frac{S_2}{(S_1)^2}. \quad (21)$$

The effective noise bandwidth ENBW is given by

$$\text{ENBW} = \text{NENBW} \cdot f_{\text{res}} = \text{NENBW} \cdot \frac{f_s}{N} = f_s \frac{S_2}{(S_1)^2}, \quad (22)$$

where f_s is the sampling frequency and f_{res} the width of one frequency bin. For the Hanning window we have $\text{NENBW} = 1.5$ bins. This equivalent noise bandwidth is required when the resulting spectrum is to be expressed as spectral density (such as for noise measurements). It can be understood by considering white noise as input to our algorithm. Due to the width of the window in the frequency domain, each frequency bin collects not only the noise in that frequency bin, but also from adjacent bins. Dividing the result by the effective noise bandwidth corrects for this phenomenon.

Another important characteristic is the maximum amplitude error e_{max} , which is the deviation of $a(f)$ from 0 dB within the range $-0.5 \leq f \leq 0.5$. It expresses the maximal possible error in the estimation of the amplitude of a sinusoidal signal that may fall anywhere within one frequency bin. For the Hanning window we find $e_{\text{max}} = -1.42$ dB. The so-called ‘flat-top’ windows (discussed in Appendix D) have a much lower e_{max} at the expense of wider bandwidth, and are hence preferable if the amplitude of sinusoidal signals must be estimated from the results.

It can be shown that the ultimate roll-off rate of the sidelobes is related to the differentiability of the window function at its boundary (where the function is extended by zero on both sides). A window function with a step discontinuity drops as $1/f$, a continuous function with a discontinuous first derivative as $1/f^2$, a window function with a continuous first derivative as $1/f^3$ etc. For the Hanning window, both the function itself and its first derivative are continuous at the boundary, and the sidelobes hence drop as $1/f^3$. For the case of window functions defined as sum of cosine terms, this relationship is further discussed in Section C.7. This ‘sidelobe dropping rate’ will be abbreviated **SLDR** in the tables below.

9 Scaling the results

Now we come back to the problem of normalizing the results of the FFT. Assume we have an input time series x_j of length N that has been converted to Volts by Equation (2). After multiplication with the chosen window function, it is subjected to a real-to-complex FFT as defined by Equation (3). We will also need the ‘window sums’ S_1 and S_2 from Equations (19) and (20). The result of the FFT is a complex vector y_m of length $N/2 + 1$ as described in Section 5. We interpret it as a power spectrum, expressed as V_{rms}^2 , as follows:

$$\text{PS}_{\text{rms}}(f_m = m \cdot f_{\text{res}}) = \frac{2 \cdot |y_m|^2}{S_1^2}; \quad m = 0 \dots N/2, \quad (23)$$

The factor S_1 takes the role of N that we had found to be necessary in Section 5. It takes into account both the length N of the DFT and the gain of the window function, plus any constant factor that was used in the computation of the window values w_j .

The factor 2 originates from the fact that we presumably use an efficient FFT algorithm that does not compute the redundant results for negative frequencies. Strictly speaking, that factor should not be applied for the first frequency bin ($m = 0$). We do, however, ignore this distinction and treat the first frequency bin identically to the other bins for the following reasons:

- The result in the first bin corresponds to zero frequency (i.e. DC), and is simply the average of the time series. If that average is of interest, the DFT methods discussed in this paper are quite unsuitable to determine it. Simple averaging or, even better, digital low-pass filtering of the time series are better procedures to determine the DC average of the time series.

- If window functions are used (as will usually be the case), there is no consistent way of normalizing the results for $m = 0$. Due to the width of the window in the frequency domain, there will be leakage of the DC average from the $m = 0$ bin into neighboring bins. In the case of flat-top windows, for example, the $m = 1$ bin will get almost the same contribution as the $m = 0$ bin. If the $m = 0$ bin were treated differently, the $m = 1$ bin would show 3 dB *more* power than the $m = 0$ bin due to the DC average, clearly an inconsistent result.
- The relationship (10) between peak amplitude and rms value does not hold for zero frequency, which further complicates the normalization problem.
- The DC average of the time series will often have been removed from the time series anyway before the FFT is performed (see Section 11).

Similarly, the last frequency bin ($m = N/2$), which corresponds to the Nyquist frequency $f_s/2$, should be treated specially in principle because it has no imaginary part if N is even. In practice this can usually be ignored also, because proper anti-aliasing filters before the A/D-converter will have removed that frequency before A/D-conversion anyway.

We now return to the scaling of the ‘normal’ results. If the desired result is a power spectral density (PSD) expressed in V^2/Hz , it is obtained by dividing the power spectrum (PS) by the effective noise-equivalent bandwidth ENBW (see Equation (22)):

$$\text{PSD}_{\text{rms}}(f_m = m \cdot f_{\text{res}}) = \frac{\text{PS}_{\text{rms}}(f_m)}{\text{ENBW}} = \frac{2 \cdot |y_m|^2}{f_s \cdot S_2}; \quad m = 0 \dots N/2, \quad (24)$$

This relationship has very important consequences. Consider the signal of the example shown in Figure 1 and let the length N of the DFT vary. The frequency resolution f_{res} and noise bandwidth ENBW both vary inversely proportional to N . In a spectrum normalized according to Equation (23), the amplitude of the peak at 1234 Hz will remain constant as expected. If we increase N and thus reduce f_{res} , the level of the noise plateau between 50 Hz and 2000 Hz, will, however, decrease in the spectrum. That is because the width of one frequency bin is reduced and consequently less and less of the equally distributed noise power falls into one frequency bin.

Computing a spectral density by normalizing according to Equation (24) compensates this effect and yields a constant level of the noise plateau independent of N , as it should be. In this normalization, however, the height of the 1234 Hz peak will increase with N .

Since the ratio between the peak and the noise plateau in the FFT output depends on N , we need to distinguish carefully between spectra and spectral densities. The magic number mentioned in Section 3 to convert them into each other is the **effective noise bandwidth ENBW, which should hence always be recorded when a spectrum or spectral density is computed**, such that the result can be converted to the other form at a later stage, when the information about the frequency resolution f_{res} and the window that was used is normally not easily available any more.

Further processing of the output is straightforward: If the desired result is a *linear spectrum* (expressed in V) or a *linear spectral density* (expressed in $V/\sqrt{\text{Hz}}$), simply take the square root of the corresponding power spectrum or spectral density:

$$\text{LSD} = \sqrt{\text{PSD}}, \quad (25)$$

$$\text{LS} = \sqrt{\text{PS}}. \quad (26)$$

Note, however, that if several spectra/spectral densities are averaged (see Section 10), this averaging must be performed with the *power* spectra/spectral densities, and the square root, if desired, must be taken only at the end.

Finally the result can be converted into other units (such as V_{pk} or dB) according to Section 6.

10 Averaging and overlap

If we compute one estimate of a spectrum with the methods described so far (i.e. multiplying one segment of the time series with a suitable window function, performing a DFT and scaling the results), we will typically find the result to be rather ‘noisy’. The theory confirms this practical observation: The standard deviation of the spectrum estimate in one frequency bin is equal to the estimate itself, i.e. 100%, if the signal in that bin is stochastic. It does not help the increase the length N of the DFT; that only reduces the width of one frequency bin without improving the variance.

The usual remedy is to take the *average* of M estimates and hence reduce the standard deviation of the averaged result by a factor of $1/\sqrt{M}$. However, the properties of the signal must remain stationary during the averaging. Note that the averaging must be done with the power spectrum (PS) or the power spectral density (PSD), not with their square roots LS or LSD. If the square roots are desired as result, they must be computed at the end after the averaging is finished. In conjunction with the use of window functions, this method of averaging several spectra is also known under names such as “Welch’s method of averaging modified periodograms”, “Welch’s overlapped Segmented average” etc.

When the spectra are averaged, the variance/standard deviation of the estimate can be computed at the same time with little extra effort. The ratio of the standard deviation to the averaged estimate itself gives some indication of the ‘randomness’ of the signal that dominates the frequency bin under consideration: It is near unity for stochastic signals such as noise, and much smaller for coherent signals such as sinusoidal peaks.

If a long continuous data stream is simply split into several non-overlapping segments of length N and each segment is processed by a DFT with a window function, we have a situation as illustrated in Figure 7.




Figure 7 *Segmented data stream with window and without overlap.*

Due to the fact that the window function is typically very small or zero near its boundaries, a significant portion of the data stream is effectively ignored in the analysis. This is clearly not optimal in those situations where the data stream was produced at great expense, and maximal possible information is to be extracted from it.³ The situation can be improved by letting the segments overlap, as illustrated in Figure 8.

³Sometimes also the opposite case occurs: when raw data are cheaply available in nearly unlimited amounts,




Figure 8 *Segmented data stream with window and overlap.*

How much should the segments overlap? That depends on the window function, and also a little on the requirements. For windows that are relatively wide in the time domain (such as Hanning), 50 % is a commonly used value for the overlap. For narrower window functions (in particular flat-top windows), a higher overlap (up to 84 %) may be appropriate.

In choosing the optimal amount of overlap, a compromise must be found between the resulting ‘flatness’ of the data weighting and the computational effort. The description of each window function in Appendix C includes a graph showing three curves that can help in choosing an appropriate overlap. Figure 9 shows these three curves for the Hanning window:




Figure 9 *Overlap characteristics of the Hanning window.*

Amplitude Flatness (AF): When several overlapping window functions are applied to the data stream, each data point is sampled several times with different weights. One possible measure of the total weight of each data point is the sum of all window values that are applied to that particular data point. Figure 10 illustrates this situation for a Hanning window with 33 % overlap (33 % is smaller than the optimum, 50 %, and was chosen in this figure for illustrative purposes). If we assume that all our data points have the same validity, we want that sum to remain fairly constant for all data points. The Amplitude Flatness (AF) is the ratio of the minimal total weight that may be applied to any data

and the computational effort to process the data is the limiting factor in the analysis. If the properties of the raw data are truly stationary, overlapping can be omitted, but it does not hurt, either, as long as the overlap correlation (see Equation (27)) remains small.

point to the maximal total weight, evaluated for one certain value of the overlap. In the example shown in Figure 10, it is 0.5. For the Hanning window, the curve of summed window values becomes exactly flat for 50 % overlap, such that Amplitude Flatness (AF) reaches its perfect value of 1.0 in this case. For most other window functions, there is no such singular point, but the amplitude flatness nevertheless approaches 1.0 very quickly with increasing amounts of overlap.




Figure 10 *Amplitude flatness of a Hanning window with 33 % overlap.*

Power Flatness (PF): While the linear summing of window values (as done in the definition of amplitude flatness above) is appropriate for the amplitude of sinusoidal signals, incoherent signals such as noise must be added quadratically. Hence we define the Power Flatness (PF) similarly as the amplitude flatness above, but with the window values from the different instances of the window function being added quadratically. The notions of Amplitude and Power Flatness were motivated by a similar discussion in [Trethewey2000].

Overlap Correlation (OC): If the overlap becomes too big, the spectrum estimates from subsequent stretches become strongly correlated, even if the signal is random. A high correlation means wasted computational effort. The Overlap Correlation (OC) for an overlap r is computed as (see [Harris78]):

$$OC(r) = \frac{\sum_{j=0}^{rN-1} w_j w_{j+(1-r)N}}{\sum_{j=0}^{N-1} w_j^2} \quad (27)$$

When choosing a value for the overlap, we will want a ‘flatness’ as high as possible, in order to give equal weight to all data. On the other hand, we want the overlap correlation to be small, such as not to waste computational effort by repeatedly computing highly correlated results. In order to get an idea of a suitable overlap, we somewhat arbitrarily define the ‘recommended overlap’ (ROV) as that value of the overlap where the distance between the amplitude flatness

(AF) and the overlap correlation (OC) reaches its maximum (see Figure 9). For the Hanning window, this maximum is reached at an overlap of 50 %. This definition⁴ happens to yield reasonable results for all the windows considered in this paper.

11 Preprocessing of the data: DC and trend removal

Often the time series we use as input for our algorithm will have a non-zero DC average, which may even slowly change over time. Such a DC average will show up in the first frequency bin ($m = 0$) of the resulting spectrum. If a window function is used or the average changes over time, it will also leak into adjacent frequency bins, possibly masking low-frequency signals.

The DC average is not usually of interest as a result of the DFT processing, due to calibration problems (see Section 9) and because there are simpler and better methods to determine it (i.e. simple averaging or digital low-pass filtering).

Hence we normally want to remove it from the data before starting the DFT processing. There are several ways to accomplish this, in (roughly) increasing order of perfection:

1. Compute the average of the whole time series (before splitting it into segments) and subtract that average from all data points.
2. Compute a straight line between the first and last data point of the whole time series (before splitting it into segments) and subtract that line from all data points.
3. Compute an average trend via linear regression of the whole time series (before splitting it into segments) and subtract that line from all data points.
4. Compute the average of each segment (before applying the window function) and subtract that average from the data points.
5. Compute a straight line between the first and last data point of each segment (before applying the window function) and subtract that line from the data points.
6. Compute an average trend via linear regression of each segment (before applying the window function) and subtract that line from the data points.
7. Pass the input time series through a digital high-pass filter.

Whether such a procedure should be applied depends on the situation. If the DC average is known to be negligible (e.g. because of AC coupling) and/or the low frequency bins are uninteresting anyway, it can be omitted. On the other hand, in some cases (e.g. small fluctuations of huge numbers such as variations in the Hz range of a beat frequency in the GHz range) it is even essential to prevent numerical collapse of the FFT algorithm.

For a general-purpose algorithm where the properties of the input data are unknown, it is recommended to apply at least a simple version of the above procedures since they are cheap compared with the rest of the processing, increase the usefulness of the lowest frequency bins and may prevent garbage results.

⁴The alternative of taking the point where the distance between PF and OC reaches a maximum yields too small overlaps for some windows.

12 Putting it all together

Here is one suggestion how all the details mentioned in this report can be put together to a working algorithm. It is, of course, not the only possibility, neither is it necessarily the best. Some of the individual steps below might be omitted, or their order might be interchanged as fits the situation at hand.

Input data: We assume to have a long stream of input data that has already been converted to floating point numbers (corresponding to ‘Volts’ at the input of the A/D converter) according to Section 4 on page 7. We also assume that the sampling frequency f_s is known and fixed.

Frequency resolution and length of DFT: Now choose a frequency resolution f_{res} . Typical values range between $f_s/100$ and $f_s/100\,000$. Find the corresponding length N of the DFT, according to $N = f_s/f_{\text{res}}$ (Equation (15)). Round N to the next convenient value (e.g. according to Equation (9) if FFTW is used), and re-compute the final frequency resolution $f_{\text{res}} = f_s/N$.

Window function: At this point a window function can be chosen and computed. The first decision to be made is whether a flat-top window is necessary. That will always be the case if the amplitude of sinusoidal peaks is to be determined from the spectrum. Otherwise, the other (i.e. non-flat-top) windows are preferable because of their smaller bandwidths. As a general rule, try to choose a window with a sidelobe suppression not less than the intrinsic signal-to-noise ratio of the input signal. If this signal-to-noise ratio is unknown, one can start with, e.g., a Hanning window and determine from the resulting spectrum how much the highest peak sticks out of the background noise. With this information, choose a suitable window from Table 2 on page 29. Good choices are the Kaiser window (see also Table 4 on page 37) and the new flat-top windows HFT xx .

Since N is already known, the window values w_j can now be computed. At the same time, compute the ‘window sums’ S_1 and S_2 (Equation (19)). If overlapping is desired, select a suitable value for the overlap, e.g. from Table 2 or the description of the selected window.

Record f_{res} , NENBW and ENBW (Equations (21) and (22)).

Splitting of the data stream: Identify continuous stretches of useful data (i.e. without glitches, pulses, missing periods etc.). Split these periods into segments of length N , overlapping as desired. Remove the DC and/or trend of each segment according to Section 11, if desired.

FFT: Multiply each segment with the pre-computed window values w_j . Put the product through the real-to-complex FFT algorithm. Unpack the result, compute the squared magnitude, and average it (separately for each frequency bin).

scaling: After all segments have been processed, scale the average according to Section 9 using Equation (23) or (24). If the linear spectrum/spectral density is required, apply Equation (25) or (26). Finally, convert to other units according to Section 6, if necessary.

13 Testing the algorithm

Once all the procedures described in the previous Section have been put together into one program, it is important to test that program. The easiest way to achieve this is to use artificial

test data. Such data should include one or several sinusoidal signals with known amplitude (yielding sharp peaks in the spectrum) as well as a noise background with known spectral density. While it is straightforward to generate a sine function, it is not so obvious how to generate noise with a known spectral density. The easiest way is to make use of the digitizing noise (further discussed in Appendix A.2 and Reference [Lyons, Section 9.3]). By rounding the simulated time-series data to integer multiples of a suitably chosen unit U_{LSB} , a noise floor with the density

$$\tilde{U}_{\text{dig}} = \frac{U_{\text{LSB}}}{\sqrt{6 \cdot f_s}} \quad (28)$$

is introduced in the spectrum. If, however, the signal consists of only one sine wave with a frequency that is harmonically related to the sampling frequency, the deviations between the ideal sine wave and its rounded samples that cause the noise floor are repetitive and hence do not have a white spectrum. By using two sine waves with at least one frequency that is not harmonically related to the sampling frequency, this effect is avoided and the noise floor can be treated as white noise. The following C program illustrates these ideas:

```
#include <stdio.h>
#include <math.h>
#define TWOPI 6.28318530717959
int main (void)
{
    double fs = 10000;           /* sampling frequency [Hz] */
    double f1 = 1234;           /* first signal frequency [Hz] */
    double amp1 = 2.82842712474619; /* 2 Vrms */
    double f2 = 2500.2157;      /* second signal frequency [Hz] */
    double amp2 = 1;            /* 0.707 Vrms */
    double ulsb = 1e-3;         /* Value of 1 LSB in Volt */
    int i;
    double t, u, ur;

    for (i = 0; i < 1000000; i++)
    {
        t = (double) i / fs;
        u = amp1 * sin (TWOPI * f1 * t) + amp2 * sin (TWOPI * f2 * t);
        ur = floor (u / ulsb + 0.5) * ulsb;      /* Rounding */
        printf ("%10.6f %8.5f\n", t, ur);        /* ASCII output */
        fwrite (&ur, sizeof (double), 1, stdout); /* alternative binary output */
    }
    return 0;
}
```

In this example, the signal consists of two sinusoidal components at 1234 Hz (amplitude $2 V_{\text{rms}} = 2.828 V_{\text{pk}}$) and 2500.2157 Hz (amplitude $0.707 V_{\text{rms}} = 1 V_{\text{pk}}$). The output samples were rounded to integer multiples of $U_{\text{LSB}} = 1 \text{ mV}$, yielding a noise background of

$$\tilde{U}_{\text{dig}} = \frac{U_{\text{LSB}}}{\sqrt{6 \cdot f_s}} = \frac{1 \text{ mV}}{\sqrt{6 \cdot 10 \text{ kHz}}} \approx 4.08 \mu\text{V}/\sqrt{\text{Hz}}. \quad (29)$$




Figure 11 Results of the analysis of the test signal generated by the example C program. Both the spectral density (left curve) and the spectrum (right curve) can be obtained with one single FFT analysis of the time series; they are related by the factor $\sqrt{ENBW} = \sqrt{12.618 \text{ Hz}} = 3.552 \sqrt{\text{Hz}}$.

Its output was analyzed by a procedure as described in Section 12. The result is shown in Figure 11, both as linear spectrum and as linear spectral density.

The data were generated for a sampling frequency of $f_s = 10 \text{ kHz}$. For the analysis, the desired frequency resolution was chosen to be $f_{\text{res}} = 3 \text{ Hz}$, resulting in a desired FFT length $N = 3333$. That number was rounded to the next convenient value according to Equation (9), yielding $N = 3328 = 13 \cdot 2^8$. This in turn gave the final frequency resolution $f_{\text{res}} = 3.0048 \text{ Hz}$. Since the amplitude of the sinusoidal signals was to be determined with high accuracy, and the dynamic range of the spectrum was a little more than 100 dB, the flat-top window HFT116D was chosen for the analysis. In order to obtain a smooth spectrum that allows an accurate determination of the noise floor, 599 spectra were averaged. The numerical results (plotted in Figure 11) show that both the amplitude of the sinusoidal components and the level of the noise background were determined correctly.

14 MATLAB pwelch function

Since many people involved in the data analysis of GEO 600 use MATLAB, which includes a function that implements Welch’s method (called ‘`pwelch`’ in the ‘Signal Processing Toolbox’), this function is briefly discussed here in relation to the conventions presented in the main text. The function in its most detailed invocation is called as

```
[Pxx,f] = pwelch(x,nwin,noverlap,nfft,fs)
```

Here $\mathbf{x}$ is a vector with the time series, $\mathbf{nwin}$ either an integer (then a Hamming window of that length is used), or a vector holding the window weights w_j , $\mathbf{noverlap}$ an integer (a fraction of the FFT length) that indicates the desired overlap, $\mathbf{nfft}$ an integer giving the length N of the FFT ($\mathbf{nfft}$ should be the same as the length of the window vector $\mathbf{nwin}$; this is not clear from the

help text), and `fs` the sampling frequency (according to the help text ‘an integer’, apparently a typo).

The results are `Pxx`, a vector with the power spectral density (PSD), and `f`, a vector with the corresponding frequencies. If y is the unit of the input time series, the output `Pxx` has the unit y^2/Hz .

If we compare the functionality of `pwelch` with what we have listed in Section 12 above, we find that

- `pwelch` nearly does what we want, except removing DC average and trend.
- If you want fancy window functions, you have to compute them yourself and pass them as vector `nwin`.
- MATLAB uses the FFTW package which means that we have the same flexibility in choosing N as given in Equation (9).
- The main deficiency of `pwelch` is that it only computes power spectral densities (PSD) and has no default way of computing power spectra (PS) or the missing factor between PSD and PS, namely the equivalent noise bandwidth (ENBW). If you compute your own window function (recommended), this can, however, easily be remedied by using Equation (22).

Note for readers in the GEO 600 project: Karsten Kötter (e-mail: ksk@aei.mpg.de) has written a MATLAB script that implements the same functionality as `pwelch` and also computes the equivalent noise bandwidth (ENBW).

Appendices

A Maximal dynamic range of real signals

Since the selection of a window function involves a trade-off between spectral leakage (i.e. how well the sidelobes are suppressed compared with the main peak) and bandwidth, it is interesting to investigate what maximal dynamic range we can expect from real signals. We consider here two common noise sources that limit the dynamic range: electronic noise and digitizing noise. The aim is to use window functions that do not deteriorate the noise floor originating from these inevitable noise sources, i.e. the spectral leakage of the window transfer function at a reasonable frequency separation should be considerably lower (e.g. at least 6 dB) than the inherent dynamic range of the input signal. This requirement is, of course, rather vague, and the answer depends on the characteristics of the particular input signal in question.

A.1 Electronic noise

All signals fed to an A/D converter must first pass through various electronic circuits which have a certain noise level. To find a lower limit for that noise, we consider as example an ultra-low-noise operational amplifier (such as AD797 or LT1028). These amplifiers typically have a useful signal range of approximately ± 12 V, and an input noise level⁵ of approximately $2 \text{ nV}/\sqrt{\text{Hz}}$. In order to compute the dynamic range we compare the maximal amplitude of a sinusoidal signal

$$U_{\text{max}} = 12 V_{\text{pk}} \approx 8.5 V_{\text{rms}}, \quad (30)$$

with the noise accumulated in one frequency bin

$$U_{\text{noise}} = 2 \frac{\text{nV}_{\text{rms}}}{\sqrt{\text{Hz}}} \cdot \sqrt{\text{ENBW}}. \quad (31)$$

For the comparison we have to specify the noise bandwidth ENBW given by

$$\text{ENBW} = \text{NENBW} \cdot f_{\text{res}}, \quad (32)$$

where f_{res} is the width of one frequency bin, and NENBW is in the range $1 \dots 4$, depending on the window function (see Equation (22)). The comparison is shown in Figure 12. If the dynamic range were limited by electronic noise alone, the requirements on the dynamic range of the window function would be rather stringent.

A.2 Digitizing noise

The noise source that usually limits the dynamic range of spectrum measurements is caused by the digitizing process introducing a noise linear spectral density given by

$$\tilde{U}_{\text{dig}} = \frac{U_{\text{LSB}}}{\sqrt{6 \cdot f_s}}, \quad (33)$$

⁵The input *voltage* noise of these amplifiers is approximately $1 \text{ nV}/\sqrt{\text{Hz}}$, but in any real application we have additional noise contributions from the thermal noise of the input impedance and the input current noise.




Figure 12 Maximal dynamic range of electronic signals, assuming a range of ± 12 V and an input noise level of $2 \text{ nV}/\sqrt{\text{Hz}}$.

where U_{LSB} is the voltage corresponding to one least significant bit and f_s the sampling frequency (see e.g. Reference [Lyons, Section 9.3]). We assume using an A/D converter with an effective resolution⁶ of n bits. The maximal amplitude of a sinusoidal signal is then given by

$$U_{\text{max, pk}} = 2^{n-1} \cdot U_{\text{LSB}}, \quad (34)$$

$$U_{\text{max, rms}} = \frac{2^{n-1}}{\sqrt{2}} U_{\text{LSB}}, \quad (35)$$

which we have to compare with the noise accumulated in one frequency bin

$$U_{\text{noise, rms}} = \tilde{U}_{\text{dig}} \cdot \sqrt{\text{ENBW}} \quad (36)$$

$$= \frac{U_{\text{LSB}}}{\sqrt{6 \cdot f_s}} \cdot \sqrt{\text{NENBW} \cdot f_{\text{res}}} \quad (37)$$

$$= U_{\text{LSB}} \cdot \sqrt{\frac{\text{NENBW}}{6N}}, \quad (38)$$

where N is the length of the DFT. The dynamic range is hence given by

$$\frac{U_{\text{max, rms}}}{U_{\text{noise, rms}}} = 2^{n-1} \cdot \sqrt{N} \cdot \sqrt{\frac{3}{\text{NENBW}}}. \quad (39)$$

This function is shown in Figure 13 for a typical value of $\text{NENBW} = 3$ bins and several values of the A/D resolution n . Since the effective resolution of modern A/D converters reaches 20 bits or better, we may expect a dynamic range of up to 180 dB under optimal conditions. In practice,

⁶The effective resolution of an A/D converter is always smaller than its nominal resolution due to imperfections and noise. On the other hand, the effective resolution can be enhanced to a value that is better than the nominal resolution by averaging several samples. Indeed one standard method to evaluate the effective resolution of an A/D converter is to digitize a known sine wave, compute a spectrum with a DFT, determine the noise floor and evaluate Equation (33).

other noise sources (such as those originating in the experiment itself, e.g. technical noise, shot noise etc.) are often likely to have a noise level higher than the digitizing noise. Hence in many cases the window functions available in the past (e.g. Hanning, Nuttall, or Salvatore's flat-top windows) have a sufficient dynamic range such as not to introduce additional noise above the noise floor that is present anyway. In favourable conditions, however, the dynamic range of those window functions is worse than the dynamic range of the rest of the signal chain, in particular in the case of flat-top windows, which had been limited to a sidelobe level of -90 dB so far. Then the new flat-top window functions introduced in this paper, having a much better dynamic range, become attractive.




Figure 13 Maximal dynamic range due to digitizing noise, computed for $NENBW=3$ bins.

B Computing the response of a window function

The response $a(f)$ of a window for a frequency offset of f bins can be computed as follows:

$$a_r(f) = \sum_{j=0}^{N-1} w_j \cos(2\pi f j/N), \quad (\text{real part}) \quad (40)$$

$$a_i(f) = \sum_{j=0}^{N-1} w_j \sin(2\pi f j/N), \quad (\text{imaginary part}) \quad (41)$$

$$a(f) = \frac{\sqrt{a_r^2 + a_i^2}}{S_1}. \quad (42)$$

Using the symmetry of the window coefficients (Equation (18)) and some trigonometric relations $a(f)$ can be computed more efficiently as follows:

$$a_0(f) = w_{N/2} + 2 \sum_{j=1}^{N/2-1} w_j \cos \left(\pi f \left[\frac{2j}{N} - 1 \right] \right), \quad (43)$$

$$a_r(f) = w_0 + \cos(\pi f) a_0(f), \quad (44)$$

$$a_i(f) = \sin(\pi f) a_0(f), \quad (45)$$

$$a(f) = \frac{\sqrt{a_r^2 + a_i^2}}{S_1}. \quad (46)$$

For the common case $w_0 = 0$ this can be further simplified:

$$a(f) = a_0(f)/S_1. \quad (47)$$

In the case of windows that are defined as a sum of m cosine terms:

$$w_j = 1 + \sum_{k=1}^m c_k \cos k \frac{2\pi j}{N}, \quad (48)$$

the transfer function can be computed directly, without fixing a finite number N of samples, and very efficiently as a sum of shifted Dirichlet kernels:

$$x(f) = \frac{1}{f} + \frac{1}{2} \sum_{k=1}^m c_k \left(\frac{1}{f+k} + \frac{1}{f-k} \right), \quad (49)$$

$$a_r(f) = \sin(2\pi f) \frac{x(f)}{2\pi}, \quad (50)$$

$$a_i(f) = (1 - \cos(2\pi f)) \frac{x(f)}{2\pi}, \quad (51)$$

$$a(f) = \sqrt{a_r^2 + a_i^2}, \quad (52)$$

for frequency offsets $f \neq 0$, and $a(0) = 1$. If f falls on an exact integer k , the above expression reduces to

$$a(f) = |c_k|/2. \quad (53)$$

Incidentally, the normalized equivalent noise bandwidth NENBW can also be computed directly for windows of the form (48):

$$\text{NENBW} = 1 + \frac{1}{2} \sum_{k=1}^m c_k^2. \quad (54)$$

C List of window functions

This appendix lists many popular and useful window functions, including several new high-performance functions recently developed by one of the authors (G.H.). Each window function is described in a text section below and graphically shown on a separate page at the end of this document. Furthermore, Table 2 summarizes the main characteristics of each window.

Name	Text Section	Graph page	PSLL [dB]	SLDR [f^{-n}]	NENBW [bins]	3 dB BW [bins]	flatness [dB]	ROV [%]
Rectangular	C.1 p.30	51	13.3	1	1.0000	0.8845	-3.9224	0.0
Welch	C.3 p.31	53	21.3	2	1.2000	1.1535	-2.2248	29.3
Bartlett	C.2 p.30	52	26.5	2	1.3333	1.2736	-1.8242	50.0
Hanning	C.4 p.31	54	31.5	3	1.5000	1.4382	-1.4236	50.0
Hamming	C.5 p.32	55	42.7	1	1.3628	1.3008	-1.7514	50.0
Nuttall3	C.7.1 p.34	57	46.7	5	1.9444	1.8496	-0.8630	64.7
Nuttall4	C.7.4 p.35	60	60.9	7	2.3100	2.1884	-0.6184	70.5
Nuttall3a	C.7.2 p.35	58	64.2	3	1.7721	1.6828	-1.0453	61.2
Kaiser3	C.8 p.36	64	69.6	1	1.7952	1.7025	-1.0226	61.9
Nuttall3b	C.7.3 p.35	59	71.5	1	1.7037	1.6162	-1.1352	59.8
Nuttall4a	C.7.5 p.35	61	82.6	5	2.1253	2.0123	-0.7321	68.0
BH92	C.6 p.32	56	92.0	1	2.0044	1.8962	-0.8256	66.1
Nuttall4b	C.7.6 p.36	62	93.3	3	2.0212	1.9122	-0.8118	66.3
Kaiser4	C.8 p.36	65	94.4	1	2.0533	1.9417	-0.7877	67.0
Nuttall4c	C.7.7 p.36	63	98.1	1	1.9761	1.8687	-0.8506	65.6
Kaiser5	C.8 p.36	66	119.8	1	2.2830	2.1553	-0.6403	70.5
Flat-top windows:								
SFT3F	D.1.1 p.40	67	31.7	3	3.1681	3.1502	+0.0082	66.7
SFT3M	D.1.4 p.41	70	44.2	1	2.9452	2.9183	-0.0115	65.5
FTNI	D.2.1 p.42	73	44.4	1	2.9656	2.9355	+0.0169	65.6
SFT4F	D.1.2 p.40	68	44.7	5	3.7970	3.7618	+0.0041	75.0
SFT5F	D.1.3 p.40	69	57.3	7	4.3412	4.2910	-0.0025	78.5
SFT4M	D.1.5 p.41	71	66.5	1	3.3868	3.3451	-0.0067	72.1
FTHP	D.2.2 p.42	74	70.4	1	3.4279	3.3846	+0.0096	72.3
HFT70	D.3.1 p.45	76	70.4	1	3.4129	3.3720	-0.0065	72.2
FTSRS	D.2.4 p.43	75	76.6	3	3.7702	3.7274	-0.0156	75.4
SFT5M	D.1.6 p.41	72	89.9	1	3.8852	3.8340	+0.0039	76.0
HFT90D	D.3.3 p.46	78	90.2	3	3.8832	3.8320	-0.0039	76.0
HFT95	D.3.2 p.46	77	95.0	1	3.8112	3.7590	+0.0044	75.6
HFT116D	D.3.4 p.47	79	116.8	3	4.2186	4.1579	-0.0028	78.2
HFT144D	D.3.5 p.47	80	144.1	3	4.5386	4.4697	+0.0021	79.9
HFT169D	D.3.6 p.47	81	169.5	3	4.8347	4.7588	+0.0017	81.2
HFT196D	D.3.7 p.48	82	196.2	3	5.1134	5.0308	+0.0013	82.3
HFT223D	D.3.8 p.48	83	223.0	3	5.3888	5.3000	-0.0011	83.3
HFT248D	D.3.9 p.49	84	248.4	3	5.6512	5.5567	+0.0009	84.1

Table 2 Summary of important characteristics of each window function discussed in this report. PSLL=‘peak sidelobe level’, SLDR=‘sidelobe drop rate’, NENBW=‘normalized equivalent noise bandwidth’, ROV=‘recommended overlap’.

C.1 Rectangular window

The rectangular window, also sometimes called ‘uniform window’, is given by

$$w_j \equiv 1, \quad (55)$$

i.e. equivalent to using no window at all. Its transfer function and characteristic bandwidths are shown in Figure 15 on page 51 and are given by

$$\text{NENBW} = 1.0000 \text{ bins}, \quad (56)$$

$$W_{3\text{dB}} = 0.8845 \text{ bins}, \quad (57)$$

$$e_{\text{max}} = -3.9224 \text{ dB} = -36.3380 \%. \quad (58)$$

The first zero is located at $f = \pm 1.00$ bins. The highest sidelobe is -13.3 dB, located at $f = \pm 1.43$ bins. The sidelobes drop at a rate of f^{-1} . Overlapping makes no sense for the Rectangular window.

While $W_{3\text{dB}}$ is the narrowest of all windows, e_{max} and the spectral leakage are the worst of all windows considered here. Although there are some special applications where the rectangular window is advantageous, it is probably not useful for any of our applications in GEO 600 and mentioned here mainly for completeness.

C.2 Bartlett window

The Bartlett window is described, e.g., in Reference [Harris78] and given by

$$z = \frac{2 \cdot j}{N}, \quad j = 0 \dots N - 1, \quad (59)$$

$$w_j = \begin{cases} z; & z \leq 1, \\ 2 - z; & z > 1, \end{cases} \quad (60)$$

i.e. a triangular function. Its transfer function and characteristics are shown in Figure 16 on page 52 and are given by

$$\text{NENBW} = 1.3333 \text{ bins}, \quad (61)$$

$$W_{3\text{dB}} = 1.2736 \text{ bins}, \quad (62)$$

$$e_{\text{max}} = -1.8242 \text{ dB} = -18.9430 \%. \quad (63)$$

The first zero is located at $f = \pm 2.00$ bins. The highest sidelobe is -26.5 dB, located at $f = \pm 2.86$ bins. The sidelobes drop at a rate of f^{-2} . At the optimal overlap of 50.0%, the amplitude flatness is 1.000, the power flatness is 0.707, and the overlap correlation is 0.250.

The Bartlett window offers a narrow bandwidth and might be useful in situations where the amplitude accuracy is unimportant and the relatively high first sidelobe does not disturb. In most situations, however, the Bartlett window does not offer any particular advantage compared with the more sophisticated windows described below and is mentioned here mainly for completeness.

C.3 Welch window

The Welch window is described, e.g., in Reference [Harris78] and given by

$$z = \frac{2 \cdot j}{N}, \quad j = 0 \dots N - 1, \quad (64)$$

$$w_j = 1 - (z - 1)^2; \quad (65)$$

i.e. a parabolic function. Its transfer function and characteristics are shown in Figure 17 on page 53 and are given by

$$\text{NENBW} = 1.2000 \text{ bins}, \quad (66)$$

$$W_{3\text{dB}} = 1.1535 \text{ bins}, \quad (67)$$

$$e_{\text{max}} = -2.2248 \text{ dB} = -22.5962 \%. \quad (68)$$

The first zero is located at $f = \pm 1.43$ bins. The highest sidelobe is -21.3 dB, located at $f = \pm 1.83$ bins. The sidelobes drop at a rate of f^{-2} . At the optimal overlap of 29.3%, the amplitude flatness is 0.828, the power flatness is 0.707, and the overlap correlation is 0.091.

The Welch window offers a narrow bandwidth and might be useful in situations where the amplitude accuracy is unimportant and the relatively high first sidelobe does not disturb. In most situations, however, the Welch window does not offer any particular advantage compared with the more sophisticated windows described below and is mentioned here mainly for completeness.

C.4 Hanning window

The Hanning window (one of a family of ‘raised cosine’ windows) is also known as ‘Hann window’. According to Reference [Harris78], it is named after the Austrian meteorologist Julius von Hann. Do not confuse it with the ‘Hamming’ window (see Section C.5). The Hanning window has already been used as example in Section 8, its properties are summarized here again for completeness. It is described, e.g., in Reference [Harris78] and given by

$$z = \frac{2\pi \cdot j}{N}, \quad j = 0 \dots N - 1, \quad (69)$$

$$w_j = \frac{1 - \cos(z)}{2} = \cos^2\left(\frac{z - \pi}{2}\right); \quad (70)$$

i.e. part of a cosine function. Its transfer function and characteristics are shown in Figure 18 on page 54 and are given by

$$\text{NENBW} = 1.5000 \text{ bins}, \quad (71)$$

$$W_{3\text{dB}} = 1.4382 \text{ bins}, \quad (72)$$

$$e_{\text{max}} = -1.4236 \text{ dB} = -15.1174 \%. \quad (73)$$

The first zero is located at $f = \pm 2.00$ bins. The highest sidelobe is -31.5 dB, located at $f = \pm 2.36$ bins. The sidelobes drop at a rate of f^{-3} . At the optimal overlap of 50.0%, the amplitude flatness is 1.000, the power flatness is 0.707, and the overlap correlation is 0.167.

The Hanning window has reasonably low spectral leakage and bandwidth and is hence used as standard window function in many commercial spectrum analyzers, if amplitude accuracy for sinusoidal signals is not important (i.e. in noise measurements).

C.5 Hamming window

The Hamming window is the simplest example of a family of windows that are constructed as a weighted sum of a constant term and some cosine terms. Do not confuse it with the ‘Hanning’ window (see Section C.4). The Hamming window is given by

$$z = \frac{2\pi \cdot j}{N}, \quad j = 0 \dots N - 1, \quad (74)$$

$$w_j = 0.54 - 0.46 \cos(z); \quad (75)$$

The coefficients are chosen such that the first sidelobe becomes nearly minimal.⁷ There is discontinuity of 0.08 at the boundary, leading to a sidelobe-drop rate of only f^{-1} . The transfer function and characteristics are shown in Figure 19 on page 55 and are given by

$$\text{NENBW} = 1.3628 \text{ bins}, \quad (76)$$

$$W_{3\text{ dB}} = 1.3008 \text{ bins}, \quad (77)$$

$$e_{\text{max}} = -1.7514 \text{ dB} = -18.2612 \%. \quad (78)$$

The first zero is located at $f = \pm 2.00$ bins. The highest sidelobe is -42.7 dB, located at $f = \pm 4.50$ bins. At the optimal overlap of 50.0%, the amplitude flatness is 1.000, the power flatness is 0.761, and the overlap correlation is 0.234. The last graph in Figure 19 on page 55 shows some discontinuities in the ‘amplitude flatness (AF)’ curve. These originate from the discontinuity at the boundary in the time domain. If, for an overlap of 50 %, two instances of the window are exactly adjacent to each other without any gap, the resulting summed amplitude becomes perfectly flat (similar to the Hanning window for 50 % overlap), and the amplitude flatness reaches unity. A gap or overlap of only one sample in the time domain, however, causes a dip or peak in the summed amplitude, which results in an amplitude flatness of less than unity. The Hamming window offers a narrow bandwidth and might be useful in situations where the amplitude accuracy is unimportant and the slow dropping rate of the sidelobes does not disturb.

C.6 Blackman-Harris window

The Blackman-Harris window is one of a family of window functions given by a sum of cosine terms. By varying the number and coefficients of the terms different characteristics can be optimized. The window we describe here is called ‘minimum 4-term Blackman-Harris window’ and was designed to have a small sidelobe adjacent to the main peak in the transfer function. Because the highest sidelobe is 92 dB below the main peak, it is also called ‘92 dB Blackman-Harris window’ to distinguish it from other members of that family. In this report its name is also abbreviated as ‘BH92’. It is described in [Harris78] and given by

$$z = \frac{2\pi \cdot j}{N}, \quad j = 0 \dots N - 1, \quad (79)$$

$$w_j = 0.35875 - 0.48829 \cos(z) + 0.14128 \cos(2z) - 0.01168 \cos(3z). \quad (80)$$

Its transfer function and characteristics are shown in Figure 20 on page 56 and are given by

$$\text{NENBW} = 2.0044 \text{ bins}, \quad (81)$$

$$W_{3\text{ dB}} = 1.8962 \text{ bins}, \quad (82)$$

$$e_{\text{max}} = -0.8256 \text{ dB} = -9.0670 \%. \quad (83)$$

⁷The minimal sidelobe level that can be reached with this form of function is -43.2 dB for a window given by $0.538379 - 0.461621 \cos(z)$.

The first zero is located at $f = \pm 4.00$ bins. The highest sidelobe is -92.0 dB, located at $f = \pm 4.52$ bins. The sidelobes drop at a rate of f^{-1} . At the optimal overlap of 66.1%, the amplitude flatness is 0.926, the power flatness is 0.718, and the overlap correlation is 0.235.

The Blackman-Harris window has very low spectral leakage combined with reasonable bandwidth and amplitude error. Due to its small sidelobes it is suitable for the detection of small sinusoidal signals adjacent in frequency to large signals. It can also be used as a general-purpose window in applications with high dynamic range if amplitude accuracy for sinusoidal signals is not very important. Note, however, that two of the Nuttall windows (called ‘Nuttall4b’ and ‘Nuttall4c’, see next Section C.7) have even lower sidelobes with the same number of cosine terms and otherwise comparable properties.

C.7 Nuttall windows

These windows are defined as a weighted sum of cosine terms (as are the Hanning, Hamming, Blackman-Harris and flat-top windows):

$$z = \frac{2\pi \cdot j}{N}, \quad j = 0 \dots N-1, \quad (84)$$

$$w_j = \sum_{k=0}^m c_k \cos(k \cdot z) = c_0 + \sum_{k=1}^m c_k \cos(k \cdot z). \quad (85)$$

Nuttall gives several sets of coefficients $\{c_k\}$ in his paper [Nuttall81], which are optimized for different aims. The coefficients are summarized in Table 3.

Name	c_0	c_1	c_2	c_3	SLDR f^{-n}	PSLL [dB]	Graph page
Nuttall3	0.375	-0.5	0.125	0	5	-46.7	57
Nuttall3a	0.40897	-0.5	0.09103	0	3	-64.2	58
Nuttall3b	0.4243801	-0.4973406	0.0782793	0	1	-71.5	59
Nuttall4	0.3125	-0.46875	0.1875	-0.03125	7	-60.9	60
Nuttall4a	0.338946	-0.481973	0.161054	-0.018027	5	-82.6	61
Nuttall4b	0.355768	-0.487396	0.144232	-0.012604	3	-93.3	62
Nuttall4c	0.3635819	-0.4891775	0.1365995	-0.0106411	1	-98.1	63

Table 3 Coefficients of Nuttall window functions, to be used with Equation (85).

In designing these window functions, a trade-off has been made between the dropping rate of the sidelobes (SLDR) and the maximal sidelobe level. After first selecting the number of terms (e.g. the window functions named Nuttall3 have 3 terms: one constant term and two cosine terms) one fixes the degree of differentiability of the window function at its boundary.

For functions of the form (85) it turns out that each extra condition on the coefficients is enough to make the function twice more differentiable, i.e. the condition

$$\sum_{k=0}^m c_k = 0 \quad (86)$$

ensures that the function is continuous at the boundary but at the same time implies that it is also differentiable once and hence has a SLDR of f^{-3} . A function that does not fulfill Equation (86) is discontinuous at the boundary and has a SLDR of f^{-1} only.

Similarly, the condition

$$\sum_{k=0}^m k^2 c_k = 0 \quad (87)$$

ensures not only twice, but also three times differentiability and hence results in a SLDR of f^{-5} , while

$$\sum_{k=0}^m k^4 c_k = 0 \quad (88)$$

yields a SLDR of f^{-7} etc.

The window functions called ‘Nuttall3’ and ‘Nuttall4’ are chosen such as to have the maximum possible degree of differentiability for a given number of terms. After fixing an arbitrary common factor⁸, there are two or three degrees of freedom available in the ‘Nuttall3’ and ‘Nuttall4’ family of windows, respectively. The subsequent functions called ‘a’, ‘b’, etc., relax the differentiability requirement by one level each and use the degrees of freedom thus becoming available to minimize the peak sidelobe level (PSLL).

The ‘Nuttall2’ window, i.e. one cosine term chosen for maximal differentiability, is just the Hanning window, while the window that might be called ‘Nuttall2a’ (minimal PSLL) would be the Hamming window.

The Nuttall windows, especially ‘Nuttall3b’, ‘Nuttall4b’ and ‘Nuttall4c’, have very low spectral leakage combined with reasonable bandwidth and amplitude error. Due to their small sidelobes they are suitable for the detection of small sinusoidal signals adjacent in frequency to large signals. They can also be used as a general-purpose window in applications with high dynamic range if amplitude accuracy for sinusoidal signals is not very important.

C.7.1 Nuttall3

The window called ‘Nuttall3’ is derived by requiring the maximal possible SLDR of f^{-5} for a three-term function. Its coefficients are taken from Ref. [Nuttall81] and listed in Table 3. The function is identical to $\cos^4(\pi(j/N - 1/2))$. The transfer function and characteristics are shown in Figure 21 and are given by

$$\text{NENBW} = 1.9444 \text{ bins}, \quad (89)$$

$$W_{3\text{dB}} = 1.8496 \text{ bins}, \quad (90)$$

$$e_{\text{max}} = -0.8630 \text{ dB} = -9.4585 \%. \quad (91)$$

The first zero is located at $f = \pm 3.00$ bins. The highest sidelobe is -46.7 dB, located at $f = \pm 3.33$ bins. At the optimal overlap of 64.7%, the amplitude flatness is 0.969, the power flatness is 0.738, and the overlap correlation is 0.228.

⁸Nuttall chose the condition $\sum (-1)^k c_k = 1$ which corresponds to normalizing the value of the window function in the center to unity: $w_{N/2} = 1$.

C.7.2 Nuttall3a

The window called ‘Nuttall3a’ is derived by requiring a SLDR of f^{-3} for a three-term function and using the remaining degree of freedom to minimize the PSL. Its coefficients are taken from Ref. [Nuttall81] and listed in Table 3. The transfer function and characteristics are shown in Figure 22 and are given by

$$\text{NENBW} = 1.7721 \text{ bins}, \quad (92)$$

$$W_{3\text{dB}} = 1.6828 \text{ bins}, \quad (93)$$

$$e_{\text{max}} = -1.0453 \text{ dB} = -11.3387 \%. \quad (94)$$

The first zero is located at $f = \pm 3.00$ bins. The highest sidelobe is -64.2 dB, located at $f = \pm 4.49$ bins. At the optimal overlap of 61.2%, the amplitude flatness is 0.943, the power flatness is 0.723, and the overlap correlation is 0.227.

C.7.3 Nuttall3b

The window called ‘Nuttall3b’ is derived by using the two degrees of freedom in a three-term function to minimize the PSL. Its coefficients are taken from Ref. [Nuttall81] and listed in Table 3. The transfer function and characteristics are shown in Figure 23 and are given by

$$\text{NENBW} = 1.7037 \text{ bins}, \quad (95)$$

$$W_{3\text{dB}} = 1.6162 \text{ bins}, \quad (96)$$

$$e_{\text{max}} = -1.1352 \text{ dB} = -12.2519 \%. \quad (97)$$

The first zero is located at $f = \pm 3.00$ bins. The highest sidelobe is -71.5 dB, located at $f = \pm 3.64$ bins. At the optimal overlap of 59.8%, the amplitude flatness is 0.939, the power flatness is 0.721, and the overlap correlation is 0.229.

C.7.4 Nuttall4

The window called ‘Nuttall4’ is derived by requiring the maximal possible SLDR of f^{-7} for a four-term function. Its coefficients are taken from Ref. [Nuttall81] and listed in Table 3. The function is identical to $\cos^6(\pi(j/N - 1/2))$. The transfer function and characteristics are shown in Figure 24 and are given by

$$\text{NENBW} = 2.3100 \text{ bins}, \quad (98)$$

$$W_{3\text{dB}} = 2.1884 \text{ bins}, \quad (99)$$

$$e_{\text{max}} = -0.6184 \text{ dB} = -6.8716 \%. \quad (100)$$

The first zero is located at $f = \pm 4.00$ bins. The highest sidelobe is -60.9 dB, located at $f = \pm 4.30$ bins. At the optimal overlap of 70.5%, the amplitude flatness is 0.937, the power flatness is 0.723, and the overlap correlation is 0.233.

C.7.5 Nuttall4a

The window called ‘Nuttall4a’ is derived by requiring a SLDR of f^{-5} for a four-term function and using the remaining degree of freedom to minimize the PSL. Its coefficients are taken from

Ref. [Nuttall81] and listed in Table 3. The transfer function and characteristics are shown in Figure 25 and are given by

$$\text{NENBW} = 2.1253 \text{ bins}, \quad (101)$$

$$W_{3\text{dB}} = 2.0123 \text{ bins}, \quad (102)$$

$$e_{\text{max}} = -0.7321 \text{ dB} = -8.0827 \%. \quad (103)$$

The first zero is located at $f = \pm 4.00$ bins. The highest sidelobe is -82.6 dB, located at $f = \pm 5.45$ bins. At the optimal overlap of 68.0%, the amplitude flatness is 0.931, the power flatness is 0.721, and the overlap correlation is 0.234.

C.7.6 Nuttall4b

The window called ‘Nuttall4b’ is derived by requiring a SLDR of f^{-3} for a four-term function and using the remaining two degrees of freedom to minimize the PSL. Its coefficients are taken from Ref. [Nuttall81] and given in Table 3. The transfer function and characteristics are shown in Figure 26 and are given by

$$\text{NENBW} = 2.0212 \text{ bins}, \quad (104)$$

$$W_{3\text{dB}} = 1.9122 \text{ bins}, \quad (105)$$

$$e_{\text{max}} = -0.8118 \text{ dB} = -8.9223 \%. \quad (106)$$

The first zero is located at $f = \pm 4.00$ bins. The highest sidelobe is -93.3 dB, located at $f = \pm 4.57$ bins. At the optimal overlap of 66.3%, the amplitude flatness is 0.924, the power flatness is 0.715, and the overlap correlation is 0.233.

C.7.7 Nuttall4c

The window called ‘Nuttall4c’ is derived by using all three degrees of freedom in a four-term function to minimize the PSL. Its coefficients are taken from Ref. [Nuttall81] and listed in Table 3. The transfer function and characteristics are shown in Figure 26 and are given by

$$\text{NENBW} = 1.9761 \text{ bins}, \quad (107)$$

$$W_{3\text{dB}} = 1.8687 \text{ bins}, \quad (108)$$

$$e_{\text{max}} = -0.8506 \text{ dB} = -9.3282 \%. \quad (109)$$

The first zero is located at $f = \pm 4.00$ bins. The highest sidelobe is -98.1 dB, located at $f = \pm 6.48$ bins. At the optimal overlap of 65.6%, the amplitude flatness is 0.923, the power flatness is 0.716, and the overlap correlation is 0.235.

C.8 Kaiser window

The Kaiser window (also called Kaiser-Bessel window) is derived by maximizing, in a certain sense, the energy that falls in the main peak of the transfer function. It is defined by (see, e.g., Reference [Harris78]):

$$z = \frac{2 \cdot j}{N} - 1, \quad j = 0 \dots N - 1, \quad (110)$$

$$w_j = \frac{I_0(\pi\alpha\sqrt{1-z^2})}{I_0(\pi\alpha)}; \quad (111)$$

α	PSLL [dB]	NENBW [bins]	3 dB BW [bins]	zero [bins]	flatness [dB]	ROV [%]
2.0	-45.9	1.4963	1.4270	2.24	-1.4527	53.4
2.5	-57.6	1.6519	1.5700	2.69	-1.2010	58.3
3.0	-69.6	1.7952	1.7025	3.16	-1.0226	61.9
3.5	-81.9	1.9284	1.8262	3.64	-0.8900	64.7
4.0	-94.4	2.0533	1.9417	4.12	-0.7877	67.0
4.5	-107.0	2.1712	2.0512	4.61	-0.7064	68.9
5.0	-119.8	2.2830	2.1553	5.10	-0.6403	70.5
5.5	-132.6	2.3898	2.2546	5.59	-0.5854	71.9
6.0	-145.5	2.4920	2.3499	6.08	-0.5392	73.1
6.5	-158.4	2.5902	2.4414	6.58	-0.4998	74.1
7.0	-171.4	2.6848	2.5297	7.07	-0.4657	75.1

Table 4 Summary of important characteristics of the Kaiser window for some values of the parameter α . PSLL=‘peak sidelobe level’, NENBW=‘normalized equivalent noise bandwidth’, ‘zero’=‘location of first zero’, ROV=‘recommended overlap’.

where I_0 is the zero-order modified Bessel function of the first kind:

$$I_0(x) = \sum_{k=0}^{\infty} \left[\frac{(x/2)^k}{k!} \right]^2, \quad (112)$$

and α is a parameter that parameterizes a trade-off between the sidelobe level and the width of the central peak. Typical values are in the range 2 to 5. The Kaiser window has a small but finite discontinuity at the boundary and hence its sidelobes drop only as f^{-1} , although the initial sidelobe level is rather low.

In order to assist in choosing appropriate values of the parameter α , Table 4 lists the important characteristics of the Kaiser windows for different values of α . We discuss here in detail and show in the Appendix F the three cases $\alpha = 3$, $\alpha = 4$ and $\alpha = 5$:

Kaiser window with $\alpha = 3$: The transfer function and characteristics are shown in Figure 28 and are given by

$$\text{NENBW} = 1.7952 \text{ bins}, \quad (113)$$

$$W_{3\text{dB}} = 1.7025 \text{ bins}, \quad (114)$$

$$e_{\text{max}} = -1.0226 \text{ dB} = -11.1069 \%. \quad (115)$$

The first zero is located at $f = \pm 3.16$ bins. The highest sidelobe is -69.6 dB, located at $f = \pm 3.32$ bins. At the optimal overlap of 61.9%, the amplitude flatness is 0.938, the power flatness is 0.722, and the overlap correlation is 0.230.

Kaiser window with $\alpha = 4$: The transfer function and characteristics are shown in Figure 29

and are given by

$$\text{NENBW} = 2.0533 \text{ bins}, \quad (116)$$

$$W_{3\text{dB}} = 1.9417 \text{ bins}, \quad (117)$$

$$e_{\text{max}} = -0.7877 \text{ dB} = -8.6697 \%. \quad (118)$$

The first zero is located at $f = \pm 4.12$ bins. The highest sidelobe is -94.4 dB, located at $f = \pm 4.25$ bins. At the optimal overlap of 67.0%, the amplitude flatness is 0.925, the power flatness is 0.719, and the overlap correlation is 0.237.

Kaiser window with $\alpha = 5$: The transfer function and characteristics are shown in Figure 30 and are given by

$$\text{NENBW} = 2.2830 \text{ bins}, \quad (119)$$

$$W_{3\text{dB}} = 2.1553 \text{ bins}, \quad (120)$$

$$e_{\text{max}} = -0.6403 \text{ dB} = -7.1061 \%. \quad (121)$$

The first zero is located at $f = \pm 5.10$ bins. The highest sidelobe is -119.8 dB, located at $f = \pm 5.20$ bins. At the optimal overlap of 70.5%, the amplitude flatness is 0.919, the power flatness is 0.717, and the overlap correlation is 0.241.

The Kaiser window is comparable in performance and potential applications to the Blackman-Harris and Nuttall windows discussed in the previous subsections. Due to its ‘tunability’ with the parameter α it can easily be optimized for a certain application. A small disadvantage might be the necessity to compute the Bessel function (112).

D Flat-Top windows

If one required result of the data processing is to determine the exact amplitude of a sinusoidal component in the input signal (as will be necessary for detector calibration purposes in GEO 600), all the windows discussed so far are rather unsatisfactory. The amplitude of the peak may be misestimated by 1 dB or more, depending on the precise location of the peak frequency within the closest frequency bin. The maximum amplitude error (as defined in Section 8.2) has been given as e_{max} in the descriptions above. For the best window so far in this respect, the ‘Nuttall4’ window, it is 0.62 dB corresponding to an amplitude error of about 7%, too much for many purposes.

Hence some people, including the manufacturers of commercial spectrum analyzers, have developed window functions that are as flat as possible in the frequency domain for frequency offsets $-0.5 \leq f \leq 0.5$ (measured in frequency bins). These are known as ‘flat-top windows’. The author is aware of only one paper, [Salvatore88], that gives formulae for flat-top windows. The windows from this paper are discussed in Section D.1.

Three more flat-top window functions that were found in various sources, such as user manuals from spectrum analyzers etc., are discussed in Section D.2.

Since none of these have a sidelobe rejection sufficient for the purposes of the GEO 600 data analysis, the author has developed several new flat-top windows, described in Section D.3 below.

All flat-top windows are defined in a manner similar to the Blackman-Harris window, i.e. as sum of a few cosine terms. All of them are *negative* in a fraction of their time-series representation.

Note also that these windows are not at all flat in the time domain (see graphs in Appendix F below).

All these flat-top windows have bandwidths of about 3...5 bins, roughly twice as wide as non-flat-top windows with comparable sidelobe suppression. In many situations this is not a severe disadvantage, because it can be compensated by choosing a smaller frequency resolution f_{res} , i.e. a larger value of the FFT length N , provided the the input data stream is long enough.

D.1 Published flat-top windows

The only paper known to the author that gives formulae of flat-top windows introduces 6 different flat-top windows (Reference [Salvatore88]).

These windows are defined as a weighted sum of cosine terms (as are the Hanning, Hamming, Blackman-Harris and Nuttall windows):

$$z = \frac{2\pi \cdot j}{N}, \quad j = 0 \dots N-1, \quad (122)$$

$$w_j = \sum_{k=0}^m c_k \cos(k \cdot z) = c_0 + \sum_{k=1}^m c_k \cos(k \cdot z). \quad (123)$$

In a manner similar to Nuttall (see Section C.7) two different window functions are defined for each (fixed) number of cosine terms by requiring different conditions that translate into equations for the unknown coefficients of the cosine terms. In addition to the differentiability conditions (86), (87) and (88), one more condition is imposed on each of the window functions below:

$$a(f_k) = 0 \quad \text{for a certain } f_k \lesssim 0.5, \quad \text{i.e. } f_k \approx 0.45. \quad (124)$$

This condition demands that the transfer function reaches a zero at a predefined frequency offset near the end of the $-0.5 \leq f \leq 0.5$ interval, thus making the transfer function nearly flat in that interval. The two types of window function for each number of terms are:

‘fast decaying’ (abbreviated ‘F’) that enforces the highest possible degree of differentiability at the boundary and thus the fastest possible sidelobe drop rate by requiring as many of the conditions (86), (87) and (88) as possible;

‘minimum sidelobe’ (abbreviated ‘M’) that uses the available degrees of freedom to minimize the level of the highest sidelobe. Actually in [Salvatore88] an analytical condition is used that enforces a zero of the transfer function at an arbitrarily selected frequency. This is not exactly equivalent to minimizing the level of the highest sidelobe, see also Section D.3.

The coefficients of the 6 window functions are summarized in Table 5. More information about each window can be found in Table 2 on page 29 and in the graph for each window, printed at the end of this report between pages 67 and 72. Note that there are some small discrepancies concerning the numbers given for the peak sidelobe level PSLR between the original paper[Salvatore88] and this work; the author believes that the numbers printed in this work are correct.

Name	c_0	c_1	c_2	c_3	c_4	SLDR f^{-n}	PSLL [dB]	Graph page
SFT3F	0.26526	-0.5	0.23474	0	0	3	-31.7	67
SFT4F	0.21706	-0.42103	0.28294	-0.07897	0	5	-44.7	68
SFT5F	0.1881	-0.36923	0.28702	-0.13077	0.02488	7	-57.3	69
SFT3M	0.28235	-0.52105	0.19659	0	0	1	-44.2	70
SFT4M	0.241906	-0.460841	0.255381	-0.041872	0	1	-66.5	71
SFT5M	0.209671	-0.407331	0.281225	-0.092669	0.0091036	1	-89.9	72

Table 5 Coefficients of Salvatore flat-top window functions, to be used with Equation (123).

D.1.1 Fast decaying 3-term flat top window

This window, called ‘SFT3F’ in this report, is derived by requiring the flatness condition (124) and a SLDR of f^{-3} . Its coefficients are taken from Ref. [Salvatore88] and listed in Table 5. The transfer function and characteristics are shown in Figure 31 on page 67 and are given by

$$\text{NENBW} = 3.1681 \text{ bins}, \quad (125)$$

$$W_{3\text{dB}} = 3.1502 \text{ bins}, \quad (126)$$

$$e_{\text{max}} = 0.0082 \text{ dB} = 0.0946 \%. \quad (127)$$

The first zero is located at $f = \pm 3.00$ bins. The highest sidelobe is -31.7 dB, located at $f = \pm 3.37$ bins. At the optimal overlap of 66.7%, the amplitude flatness is 0.998, the power flatness is 0.558, and the overlap correlation is -0.029.

D.1.2 Fast decaying 4-term flat top window

This window, called ‘SFT4F’ in this report, is derived by requiring the flatness condition (124) and a SLDR of f^{-5} . Its coefficients are taken from Ref. [Salvatore88] and listed in Table 5. The transfer function and characteristics are shown in Figure 32 on page 68 and are given by

$$\text{NENBW} = 3.7970 \text{ bins}, \quad (128)$$

$$W_{3\text{dB}} = 3.7618 \text{ bins}, \quad (129)$$

$$e_{\text{max}} = 0.0041 \text{ dB} = 0.0472 \%. \quad (130)$$

The first zero is located at $f = \pm 4.00$ bins. The highest sidelobe is -44.7 dB, located at $f = \pm 4.33$ bins. At the optimal overlap of 75.0%, the amplitude flatness is 1.000, the power flatness is 0.647, and the overlap correlation is 0.039.

D.1.3 Fast decaying 5-term flat top window

This window, called ‘SFT5F’ in this report, is derived by requiring the flatness condition (124) and a SLDR of f^{-7} . Its coefficients are taken from Ref. [Salvatore88] and listed in Table 5. The

transfer function and characteristics are shown in Figure 33 on page 69 and are given by

$$\text{NENBW} = 4.3412 \text{ bins}, \quad (131)$$

$$W_{3\text{dB}} = 4.2910 \text{ bins}, \quad (132)$$

$$e_{\text{max}} = -0.0025 \text{ dB} = -0.0282 \%. \quad (133)$$

The first zero is located at $f = \pm 5.00$ bins. The highest sidelobe is -57.3 dB, located at $f = \pm 5.31$ bins. At the optimal overlap of 78.5%, the amplitude flatness is 0.969, the power flatness is 0.648, and the overlap correlation is 0.052.

D.1.4 Minimum sidelobe 3-term flat top window

This window, called ‘SFT3M’ in this report, is derived by requiring the flatness condition (124) and using the remaining degree of freedom to reduce the peak sidelobe level. Its coefficients are taken from Ref. [Salvatore88] and listed in Table 5. The transfer function and characteristics are shown in Figure 34 on page 70 and are given by

$$\text{NENBW} = 2.9452 \text{ bins}, \quad (134)$$

$$W_{3\text{dB}} = 2.9183 \text{ bins}, \quad (135)$$

$$e_{\text{max}} = -0.0115 \text{ dB} = -0.1323 \%. \quad (136)$$

The first zero is located at $f = \pm 3.00$ bins. The highest sidelobe is -44.2 dB, located at $f = \pm 5.50$ bins. The sidelobes drop at a rate of f^{-1} . At the optimal overlap of 65.5%, the amplitude flatness is 0.949, the power flatness is 0.584, and the overlap correlation is -0.005.

D.1.5 Minimum sidelobe 4-term flat top window

This window, called ‘SFT4M’ in this report, is derived by requiring the flatness condition (124) and using the remaining two degrees of freedom to reduce the peak sidelobe level. Its coefficients are taken from Ref. [Salvatore88] and listed in Table 5. The transfer function and characteristics are shown in Figure 35 on page 71 and are given by

$$\text{NENBW} = 3.3868 \text{ bins}, \quad (137)$$

$$W_{3\text{dB}} = 3.3451 \text{ bins}, \quad (138)$$

$$e_{\text{max}} = -0.0067 \text{ dB} = 0.0776 \%. \quad (139)$$

The first zero is located at $f = \pm 4.00$ bins. The highest sidelobe is -66.5 dB, located at $f = \pm 10.50$ bins. The sidelobes drop at a rate of f^{-1} . At the optimal overlap of 72.1%, the amplitude flatness is 0.964, the power flatness is 0.641, and the overlap correlation is 0.044.

The HFT70 window (see Section D.3.1 on page 45) was designed with the same aims and reaches a better sidelobe suppression due to a better optimization procedure.

D.1.6 Minimum sidelobe 5-term flat top window

This window, called ‘SFT5M’ in this report, is derived by requiring the flatness condition (124) and using the remaining two degrees of freedom to reduce the peak sidelobe level.⁹ Its coefficients

⁹In the original work [Salvatore88] it is claimed that the remaining degree of freedom was used to enforce condition (86) and thus ensure a SLDR of f^{-3} . The coefficients listed in the paper do not, however, exactly fulfill this condition and the sidelobes accordingly drop only as f^{-1} for $f > 300$, see Figure 36. The window HFT90D (see Section D.3.3 and Figure 42 on page 78) is very similar to SFT5M and exactly fulfills condition (86).

are taken from Ref. [Salvatore88] and listed in Table 5. The transfer function and characteristics are shown in Figure 35 on page 71 and are given by

$$\text{NENBW} = 3.8852 \text{ bins}, \quad (140)$$

$$W_{3\text{dB}} = 3.8340 \text{ bins}, \quad (141)$$

$$e_{\text{max}} = 0.0039 \text{ dB} = 0.0449 \%. \quad (142)$$

The first zero is located at $f = \pm 5.00$ bins. The highest sidelobe is -89.9 dB, located at $f = \pm 5.12$ bins. The sidelobes drop at a rate of f^{-1} . At the optimal overlap of 76.0%, the amplitude flatness is 0.953, the power flatness is 0.645, and the overlap correlation is 0.053.

D.2 Flat-top windows by instrument makers

Because of the practical importance of flat-top windows and the lack (until recently) of published information, the makers of spectrum analyzers and related software have developed their own flat-top windows. All those that the author could find in user manuals, web pages, usenet discussions etc. are discussed in this Section.

D.2.1 National Instruments flat-top window

This window function, abbreviated ‘FTNI’ in this report, is cited by National Instruments on its web page and is defined by

$$z = \frac{2\pi \cdot j}{N}, \quad j = 0 \dots N-1, \quad (143)$$

$$w_j = 0.2810639 - 0.5208972 \cos(z) + 0.1980399 \cos(2z). \quad (144)$$

The transfer function and characteristics are shown in Figure 37 on page 73 and are given by

$$\text{NENBW} = 2.9656 \text{ bins}, \quad (145)$$

$$W_{3\text{dB}} = 2.9355 \text{ bins}, \quad (146)$$

$$e_{\text{max}} = 0.0169 \text{ dB} = 0.1946 \%. \quad (147)$$

The first zero is located at $f = \pm 3.00$ bins. The highest sidelobe is -44.4 dB, located at $f = \pm 3.23$ bins. The sidelobes drop at a rate of f^{-1} . At the optimal overlap of 65.6%, the amplitude flatness is 0.950, the power flatness is 0.584, and the overlap correlation is -0.007. This window is very similar to SFT3M (see Section D.1.4 and Figure 34 on page 70).

D.2.2 Old HP flat-top window

This window function, abbreviated ‘FTHP’ in this report, had been used in some older HP spectrum analyzers and is defined by

$$z = \frac{2\pi \cdot j}{N}, \quad j = 0 \dots N-1, \quad (148)$$

$$w_j = 1.0 - 1.912510941 \cos(z) + 1.079173272 \cos(2z) - 0.1832630879 \cos(3z). \quad (149)$$

Its transfer function and characteristics are shown in Figure 38 on page 74 and are given by

$$\text{NENBW} = 3.4279 \text{ bins}, \quad (150)$$

$$W_{3\text{dB}} = 3.3846 \text{ bins}, \quad (151)$$

$$e_{\text{max}} = 0.0096 \text{ dB} = 0.1103 \%. \quad (152)$$

The first zero is located at $f = \pm 4.00$ bins. The highest sidelobe is -70.4 dB, located at $f = \pm 4.65$ bins. The sidelobes drop at a rate of f^{-1} . At the optimal overlap of 72.3%, the amplitude flatness is 0.966, the power flatness is 0.640, and the overlap correlation is 0.041.

The HFT70 window (see Section D.3.1) is very similar to the FTHP window, it reaches the same sidelobe level with a smaller amplitude error e_{max} and a marginally smaller bandwidth.

D.2.3 Undisclosed HP flat-top window

Newer spectrum analyzers from HP/Agilent (such as the 3562 and 35670A) use another flat-top window. Agilent insists that the precise definition is ‘proprietary’ and cannot be given out without a non-disclosure agreement, disclosing only that it is ‘a basic cosine function with 4 terms’. It is, however, possible to measure the transfer function of the flat-top window and compare it with the known flat-top windows. It turns out that the closest match is the ‘HFT95’ window (see Section D.3.2 and Figure 41 on page 77). The comparison is shown in Figure 14.

It can be seen that the 35670A flat-top window is very close to the ‘HFT95’ window, which also has a similar equivalent noise bandwidth. That value could not be found in the documentation either, but can be deduced from the conversion factor between measurements in V and $V/\sqrt{\text{Hz}}$, as described in Section 9. The normalized equivalent noise bandwidth of the 35670A flat-top window, as programmed into its software, is

$$\text{NENBW} = 3.819 \text{ bins}, \quad (153)$$

compared with $\text{NENBW} = 3.811$ bins for the ‘HFT95’ window.

D.2.4 Stanford Research flat-top window

This window function, abbreviated ‘FTSRS’ in this report, is used in the Stanford Research SR785 spectrum analyzer and is defined by (according to the user manual):

$$z = \frac{2\pi \cdot j}{N}, \quad j = 0 \dots N-1, \quad (154)$$

$$w_j = 1.0 - 1.93 \cos(z) + 1.29 \cos(2z) - 0.388 \cos(3z) + 0.028 \cos(4z). \quad (155)$$

Its transfer function and characteristics are shown in Figure 39 on page 75 and are given by

$$\text{NENBW} = 3.7702 \text{ bins}, \quad (156)$$

$$W_{3\text{dB}} = 3.7274 \text{ bins}, \quad (157)$$

$$e_{\text{max}} = -0.0156 \text{ dB} = -0.1796 \%. \quad (158)$$

The first zero is located at $f = \pm 4.72$ bins. The highest sidelobe is -76.6 dB, located at $f = \pm 5.37$ bins. The sidelobes drop at a rate of f^{-3} . At the optimal overlap of 75.4%, the amplitude flatness is 0.958, the power flatness is 0.647, and the overlap correlation is 0.055.




Figure 14 Measured flat-top window of Agilent 35670A dynamic signal analyzer. The black asterisks represent the measured data, while the solid curve show the computed transfer function of the ‘HFT95’ flat-top window discussed in the text. The analyzer was set to a resolution of 1600 lines. For frequency offsets more than approximately ± 5 bins the measurement was dominated by the analyzer noise floor, which was approximately -93 dB in this measurement.

D.3 New flat-top windows

Since the flat-top windows available so far are not sufficient for the the data analysis of GEO 600 in all situations,¹⁰ new flat-top windows with improved characteristics have been developed by the author. Their basic form is again a sum of cosine terms, with the first term c_0 arbitrarily set to unity:¹¹

$$z = \frac{2\pi \cdot j}{N}, \quad j = 0 \dots N-1, \quad (159)$$

$$w_j = 1 + \sum_{k=1}^m c_k \cos(k \cdot z). \quad (160)$$

Most of the functions presented below additionally fulfill the condition

$$1 + \sum_{k=1}^m c_k = 0 \quad (161)$$

which ensures that the function is continuous at the boundary but at the same time implies that is also differentiable once and hence has a sidelobe dropping rate (SLDR) of f^{-3} . These functions are called ‘HFT xx **D**’ in this report, where the ‘**D**’ stands for ‘differentiable’ and xx is the maximal sidelobe level in dB.¹² Those functions that do not fulfill Equation (161) are discontinuous at the boundary, have a SLDR of f^{-1} only and are called ‘HFT xx ’

The functions were developed by applying nonlinear optimization algorithms to a ‘figure-of-merit’ function that was formed as a combination of the flatness for $-0.5 \leq f \leq 0.5$, the width of the main lobe, and the peak sidelobe level. The parameters to be varied by the algorithms were the coefficients c_k of Equation (160). Up to 8 parameters were used in the optimization. The resulting window functions are presented below, with a short text section giving their formula and describing their main characteristics, and a detailed graph near the end of this report.

D.3.1 HFT70

This window was optimized for the lowest sidelobe level that is achievable with 3 cosine terms, i.e. the same aims as were used for the SFT4M window (see Section D.1.5). Because of the better optimization procedure, a sidelobe level of -70.4 dB was reached instead of -66.5 dB in the case of the SFT4M window. The window is defined as (see also Equation (160)):

$$w_j = 1 - 1.90796 \cos(z) + 1.07349 \cos(2z) - 0.18199 \cos(3z). \quad (162)$$

Its transfer function and characteristics are shown in Figure 40 on page 76 and are given by

$$\text{NENBW} = 3.4129 \text{ bins}, \quad (163)$$

$$W_{3\text{dB}} = 3.3720 \text{ bins}, \quad (164)$$

$$e_{\text{max}} = -0.0065 \text{ dB} = 0.0750 \%. \quad (165)$$

¹⁰The dynamic range of signals digitized with the main data acquisition system, for example, can reach more than 120 dB, while the best flat-top window known so far has a sidelobe level of 90 dB only.

¹¹The window can be defined with an arbitrary factor due to our use of the ‘window sums’ S_1 and S_2 in the normalization of the results (see Equations (23) and (24)). Setting the first term c_0 to unity simplified the optimization procedure that was used to find the coefficients c_k . Another possible normalization is the one chosen by Nuttall (see the footnote in Section C.7)

¹²The ‘H’ can stand for ‘high dynamic range’, ‘high sidelobe suppression’ etc.

The first zero is located at $f = \pm 4.00$ bins. The highest sidelobe is -70.4 dB, located at $f = \pm 4.65$ bins. The sidelobes drop at a rate of f^{-1} . At the optimal overlap of 72.2%, the amplitude flatness is 0.964, the power flatness is 0.637, and the overlap correlation is 0.041.

D.3.2 HFT95

This window was optimized for the lowest sidelobe level that is achievable with 4 cosine terms. Its characteristics are very similar to those of the flat-top window that is used in newer HP/Agilent spectrum analyzers, see Section D.2.3. The window is defined as (see also Equation (160)):

$$w_j = 1 - 1.9383379 \cos(z) + 1.3045202 \cos(2z) - 0.4028270 \cos(3z) + 0.0350665 \cos(4z). \quad (166)$$

Its transfer function and characteristics are shown in Figure 41 on page 77 and are given by

$$\text{NENBW} = 3.8112 \text{ bins}, \quad (167)$$

$$W_{3\text{dB}} = 3.7590 \text{ bins}, \quad (168)$$

$$e_{\max} = 0.0044 \text{ dB} = 0.0507 \%. \quad (169)$$

The first zero is located at $f = \pm 5.00$ bins. The highest sidelobe is -95.0 dB, located at $f = \pm 7.49$ bins. The sidelobes drop at a rate of f^{-1} . At the optimal overlap of 75.6%, the amplitude flatness is 0.952, the power flatness is 0.647, and the overlap correlation is 0.056.

D.3.3 HFT90D

This window was optimized for the lowest sidelobe level that is achievable with 4 cosine terms if condition (161) is additionally imposed to ensure a sidelobe-drop rate of f^{-3} . It is very similar to the SFT5M window (see Section D.1.6), although HFT90D achieves the same sidelobe level at a marginally smaller bandwidth and with a sidelobe drop rate of f^{-3} instead of f^{-1} . The window is defined as (see also Equation (160)):

$$w_j = 1 - 1.942604 \cos(z) + 1.340318 \cos(2z) - 0.440811 \cos(3z) + 0.043097 \cos(4z). \quad (170)$$

Its transfer function and characteristics are shown in Figure 42 on page 78 and are given by

$$\text{NENBW} = 3.8832 \text{ bins}, \quad (171)$$

$$W_{3\text{dB}} = 3.8320 \text{ bins}, \quad (172)$$

$$e_{\max} = -0.0039 \text{ dB} = 0.0450 \%. \quad (173)$$

The first zero is located at $f = \pm 5.00$ bins. The highest sidelobe is -90.2 dB, located at $f = \pm 5.58$ bins. The sidelobes drop at a rate of f^{-3} . At the optimal overlap of 76.0%, the amplitude flatness is 0.953, the power flatness is 0.646, and the overlap correlation is 0.054.

D.3.4 HFT116D

This window was optimized for the lowest sidelobe level that is achievable with 5 cosine terms if condition (161) is additionally imposed to ensure a sidelobe-drop rate of f^{-3} . The window is defined as (see also Equation (160)):

$$w_j = 1 - 1.9575375 \cos(z) + 1.4780705 \cos(2z) - 0.6367431 \cos(3z) + 0.1228389 \cos(4z) - 0.0066288 \cos(5z). \quad (174)$$

Its transfer function and characteristics are shown in Figure 43 on page 79 and are given by

$$\text{NENBW} = 4.2186 \text{ bins}, \quad (175)$$

$$W_{3\text{dB}} = 4.1579 \text{ bins}, \quad (176)$$

$$e_{\max} = -0.0028 \text{ dB} = 0.0326 \%. \quad (177)$$

The first zero is located at $f = \pm 6.00$ bins. The highest sidelobe is -116.8 dB, located at $f = \pm 7.52$ bins. At the optimal overlap of 78.2%, the amplitude flatness is 0.947, the power flatness is 0.651, and the overlap correlation is 0.063.

D.3.5 HFT144D

This window was optimized for the lowest sidelobe level that is achievable with 6 cosine terms if condition (161) is additionally imposed to ensure a sidelobe-drop rate of f^{-3} . The window is defined as (see also Equation (160)):

$$w_j = 1 - 1.96760033 \cos(z) + 1.57983607 \cos(2z) - 0.81123644 \cos(3z) + 0.22583558 \cos(4z) - 0.02773848 \cos(5z) + 0.00090360 \cos(6z). \quad (178)$$

Its transfer function and characteristics are shown in Figure 44 on page 80 and are given by

$$\text{NENBW} = 4.5386 \text{ bins}, \quad (179)$$

$$W_{3\text{dB}} = 4.4697 \text{ bins}, \quad (180)$$

$$e_{\max} = 0.0021 \text{ dB} = 0.0245 \%. \quad (181)$$

The first zero is located at $f = \pm 7.00$ bins. The highest sidelobe is -144.1 dB, located at $f = \pm 7.07$ bins. At the optimal overlap of 79.9%, the amplitude flatness is 0.942, the power flatness is 0.655, and the overlap correlation is 0.069.

D.3.6 HFT169D

This window was optimized for the lowest sidelobe level that is achievable with 7 cosine terms if condition (161) is additionally imposed to ensure a sidelobe-drop rate of f^{-3} . The window is defined as (see also Equation (160)):

$$w_j = 1 - 1.97441842 \cos(z) + 1.65409888 \cos(2z) - 0.95788186 \cos(3z) + 0.33673420 \cos(4z) - 0.06364621 \cos(5z) + 0.00521942 \cos(6z) - 0.00010599 \cos(7z). \quad (182)$$

Its transfer function and characteristics are shown in Figure 45 on page 81 and are given by

$$\text{NENBW} = 4.8347 \text{ bins}, \quad (183)$$

$$W_{3\text{ dB}} = 4.7588 \text{ bins}, \quad (184)$$

$$e_{\text{max}} = 0.0017 \text{ dB} = 0.0191 \%. \quad (185)$$

The first zero is located at $f = \pm 8.00$ bins. The highest sidelobe is -169.5 dB, located at $f = \pm 10.41$ bins. At the optimal overlap of 81.2%, the amplitude flatness is 0.938, the power flatness is 0.654, and the overlap correlation is 0.072.

D.3.7 HFT196D

This window was optimized for the lowest sidelobe level that is achievable with 8 cosine terms if condition (161) is additionally imposed to ensure a sidelobe-drop rate of f^{-3} . The window is defined as (see also Equation (160)):

$$\begin{aligned} w_j = & 1 - 1.979280420 \cos(z) + 1.710288951 \cos(2z) - 1.081629853 \cos(3z) \\ & + 0.448734314 \cos(4z) - 0.112376628 \cos(5z) + 0.015122992 \cos(6z) \\ & - 0.000871252 \cos(7z) + 0.000011896 \cos(8z). \end{aligned} \quad (186)$$

Its transfer function and characteristics are shown in Figure 46 on page 82 and are given by

$$\text{NENBW} = 5.1134 \text{ bins}, \quad (187)$$

$$W_{3\text{ dB}} = 5.0308 \text{ bins}, \quad (188)$$

$$e_{\text{max}} = 0.0013 \text{ dB} = 0.0153 \%. \quad (189)$$

The first zero is located at $f = \pm 9.00$ bins. The highest sidelobe is -196.2 dB, located at $f = \pm 9.06$ bins. At the optimal overlap of 82.3%, the amplitude flatness is 0.936, the power flatness is 0.656, and the overlap correlation is 0.075.

D.3.8 HFT223D

This window was optimized for the lowest sidelobe level that is achievable with 9 cosine terms if condition (161) is additionally imposed to ensure a sidelobe-drop rate of f^{-3} . The window is defined as (see also Equation (160)):

$$\begin{aligned} w_j = & 1 - 1.98298997309 \cos(z) + 1.75556083063 \cos(2z) \\ & - 1.19037717712 \cos(3z) + 0.56155440797 \cos(4z) \\ & - 0.17296769663 \cos(5z) + 0.03233247087 \cos(6z) \\ & - 0.00324954578 \cos(7z) + 0.00013801040 \cos(8z) \\ & - 0.00000132725 \cos(9z). \end{aligned} \quad (190)$$

Its transfer function and characteristics are shown in Figure 47 on page 83 and are given by

$$\text{NENBW} = 5.3888 \text{ bins}, \quad (191)$$

$$W_{3\text{ dB}} = 5.3000 \text{ bins}, \quad (192)$$

$$e_{\text{max}} = -0.0011 \text{ dB} = -0.0129 \%. \quad (193)$$

The first zero is located at $f = \pm 10.00$ bins. The highest sidelobe is -223.0 dB, located at $f = \pm 11.38$ bins. At the optimal overlap of 83.3%, the amplitude flatness is 0.936, the power flatness is 0.659, and the overlap correlation is 0.079.

D.3.9 HFT248D

This window was optimized for the lowest sidelobe level that is achievable with 10 cosine terms if condition (161) is additionally imposed to ensure a sidelobe-drop rate of f^{-3} . The window is defined as (see also Equation (160)):

$$\begin{aligned} w_j = & 1 - 1.985844164102 \cos(z) + 1.791176438506 \cos(2z) \\ & - 1.282075284005 \cos(3z) + 0.667777530266 \cos(4z) \\ & - 0.240160796576 \cos(5z) + 0.056656381764 \cos(6z) \\ & - 0.008134974479 \cos(7z) + 0.000624544650 \cos(8z) \\ & - 0.000019808998 \cos(9z) + 0.000000132974 \cos(10z). \end{aligned} \quad (194)$$

Its transfer function and characteristics are shown in Figure 48 on page 84 and are given by

$$\text{NENBW} = 5.6512 \text{ bins}, \quad (195)$$

$$W_{3\text{dB}} = 5.5567 \text{ bins}, \quad (196)$$

$$e_{\max} = 0.0009 \text{ dB} = 0.0104 \%. \quad (197)$$

The first zero is located at $f = \pm 11.00$ bins. The highest sidelobe is -248.4 dB, located at $f = \pm 13.37$ bins. At the optimal overlap of 84.1%, the amplitude flatness is 0.934, the power flatness is 0.659, and the overlap correlation is 0.080.

E Bibliography and Acknowledgements

References

- [FFTW] FFTW subroutine package, available under GPL from
<http://www.fftw.org/>
- [Harris78] Fredric J. Harris: ‘On the Use of Windows for Harmonic Analysis with the Discrete Fourier Transform’, Proceedings of the IEEE, Vol. 66 (1): 51–83, Jan 1978
- [Nuttall81] A.H. Nuttall: ‘Some windows with very good sidelobe behaviour’, IEEE Trans. on Acoustics, Speech and Signal Processing, Vol ASSP-29 (1): 84–91 (1981).
- [Salvatore88] L. Salvatore, A. Trotta: ‘Flat-top windows for PWM waveform processing via DFT’, IEE proceedings, Vol. 135, Part B, (6) 346–361 (1988).
- [Trethewey2000] M.W. Trethewey: ‘Window and overlap processing effects on power estimates from spectra’, Mech. Syst. and Signal Process., Vol. 14 (2) 267–278, 2000.
- [Lyons] Richard G. Lyons: ‘Understanding Digital Signal Processing’, Addison-Wesley 1997.
- [Press] W.H. Press, S.A. Teukolsky, W.T. Vetterling, B.P. Flannery: ‘Numerical Recipes in C, The Art of Scientific Computing’, Cambridge University Press, 2nd ed., 1992.
- [Schutz] B.F. Schutz: ‘Relation between Continuous and Discrete Fourier Transforms’, internal note, April 1992.
- [1] 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 Electroacoustics, Vol. AU-15 (June 1967), 70-73.

Acknowledgements

The authors wish to thank K. Kawabe, K. Kötter, A. Weidner and U. Weiland for their proof-reading of the manuscript and their various useful comments.

F Graphs of window functions

This appendix shows one comprehensive graph that includes the most important characteristics for each window function discussed in the text. The first plot for each window shows the window function in the time domain. For the plot it is normalized by computing

$$w'_j = w_j \frac{N}{S_1}. \quad (198)$$

Such a normalization is not necessary for the real application of the window, since we will use S_1 in the scaling of the results (see Section 9). In fact, we can compute the window values w_j with any convenient constant factor for this reason.

The next four plots show the transfer function of the window function in the frequency domain, with different x -axes. In the second of these plots the 3 dB bandwidth $W_{3\text{dB}}$ is indicated with two small red arrows. The last of these plots has a logarithmic x -axis up to 10 000 bins in order to show the decay of the sidelobes far away from the main peak. For some of the window functions which have a very steep decay (see e.g. Figure 24 on page 60), the sidelobe level falls below -300 dB or so within the plot range and cannot be computed numerically any more; in these cases the last part of the curve shown is dominated by numerical noise. When comparing window functions, note that the y -axes are different for each window.

The last plot shows the overlap behaviour of the window, as discussed in Section 10. The vertical arrow is drawn at the position of the recommended overlap (ROV), as defined in that Section. For those windows that do not drop to zero at their end, again there is some numerical instability in the computation which shows up as non-smooth steps mainly in the ‘amplitude flatness’ curve (see also the discussion in Section C.5 on page 32 about the Hamming window, where this effect is most noticeable). For the flat-top windows, which are negative in a fraction of their time-domain representation, the amplitude flatness (AF) and overlap correlation (OC) also become negative for small values of the overlap.

All graphs and bandwidths below have been numerically computed with actual discrete window functions, using $N = 1000$, except for the 5th plot (logarithmic x -axis up to 10 000 bins), which uses $N = 30\,000$.

Rectangular




Figure 15 Rectangular window. See text Section C.1 on page 30.

Bartlett




Figure 16 *Bartlett window. See text Section C.2 on page 30.*

Welch




Figure 17 Welch window. See text Section C.3 on page 31.

Hanning




Figure 18 Hanning window. See text Section C.4 on page 31.

Hamming




Figure 19 *Hamming window. See text Section C.5 on page 32.*

BH92




Figure 20 92 dB Blackman-Harris window. See text Section C.6 on page 32.

Nuttall3




Figure 21 Nuttall3 window. See text Section C.7.1 on page 34.

Nuttall3a




Figure 22 Nuttall3a window. See text Section C.7.2 on page 35.

Nuttall3b




Figure 23 Nuttall3b window. See text Section C.7.3 on page 35.

Nuttall4




Figure 24 Nuttall4 window. See text Section C.7.4 on page 35.

Nuttall4a




Figure 25 Nuttall4a window. See text Section C.7.5 on page 35.

Nuttall4b




Figure 26 Nuttall4b window. See text Section C.7.6 on page 36.

Nuttall4c




Figure 27 Nuttall4c window. See text Section C.7.7 on page 36.

Kaiser3




Figure 28 Kaiser window with $\alpha = 3$. See text Section C.8 on page 36.

Kaiser4




Figure 29 Kaiser window with $\alpha = 4$. See text Section C.8 on page 36.

Kaiser5




Figure 30 Kaiser window with $\alpha = 5$. See text Section C.8 on page 36.

SFT3F




Figure 31 Fast decaying 3-term flat top window. See text Section D.1 on page 39.

SFT4F




Figure 32 Fast decaying 4-term flat top window. See text Section D.1 on page 39.

SFT5F




Figure 33 Fast decaying 5-term flat top window. See text Section D.1 on page 39.

SFT3M




Figure 34 Minimum sidelobe 3-term flat top window. See text Section D.1 on page 39.

SFT4M




Figure 35 Minimum sidelobe 4-term flat top window. See text Section D.1 on page 39.

SFT5M




Figure 36 Minimum sidelobe 5-term flat top window. See text Section D.1 on page 39.

FTNI




Figure 37 Flat top window as described by National Instruments. See text Section D.2.1 on page 42.

FTHP




Figure 38 Flat top window used in older HP spectrum analyzers. See text Section D.2.2 on page 42.

FTSRS




Figure 39 Flat top window used in the Stanford SR785 spectrum analyzer. See text Section D.2.4 on page 43.

HFT70




Figure 40 HFT70 Flat top window. See text Section D.3.1 on page 45.

HFT95




Figure 41 HFT95 Flat top window. See text Section D.3.2 on page 46.

HFT90D




Figure 42 HFT90D Flat top window. See text Section D.3.3 on page 46.

HFT116D




Figure 43 HFT116D Flat top window. See text Section D.3.4 on page 47.

HFT144D




Figure 44 *HFT144D Flat top window. See text Section D.3.5 on page 47.*

HFT169D




Figure 45 HFT169D Flat top window. See text Section D.3.6 on page 47.

HFT196D




Figure 46 HFT196D Flat top window. See text Section D.3.7 on page 48.

HFT223D




Figure 47 HFT223D Flat top window. See text Section D.3.8 on page 48.

HFT248D




Figure 48 HFT248D Flat top window. See text Section D.3.9 on page 49.