Advanced AD/DA converters

FFT, ΔΣ modulators and windowing

Pietro Andreani
Dept. of Electrical and Information Technology
Lund University, Sweden

Overview

- Principles of FFT
- Coherent and incoherent sampling
- Noise folding
- Windowing
- FFT scaling
- Noise bandwidth

About the FFT – I

The FFT $X(f)$ of a sample sequence $x(n)$ of length N is made of N discrete lines (***bins***) at frequencies 0, 1/N, 2/N, …, N-1/N, and is defined (assuming a normalized sample rate of 1Hz) as

$$X(f)=\sum_{n=0}^{N-1} x(n)\cdot e^{-jn2\pi f}$$

The spectrum is symmetrical around the Nyquist frequency = 1/2 → only half of the FFT computations are usually represented → N/2 spectral lines in the base band

Each line gives the spectral power falling in a band of width 1/N and centered around the line itself

Thus, the FFT operates like a spectrum analyzer with N channels each with bandwidth 1/N

If the number of points in the series increases, the channel bandwidth of the equivalent spectrum analyzer decreases and each channel will contain less noise power

---

FFT – processing gain

The FFT of the quantization noise is, on average, a factor $\frac{3}{2}\cdot 2^{2M}\cdot\frac{N}{2}$ below the full scale, where **M** is the number of bits in the quantizer, and **N** is the number of data samples. The overall noise floor becomes

$$\left.x_{noise}^{2}\right|_{dB} = P_{sig} - 1.76 - 6.02\cdot M - 10\cdot\log\frac{N}{2}$$

The term $10\cdot\log\frac{N}{2}$ is called *processing gain* of the FFT

Periodicity assumption in FFT

FFT assumes the input to be N-periodic. Real signals are never periodic, and the N-periodicity assumption leads to discontinuities between the last and first samples of successive sequences.

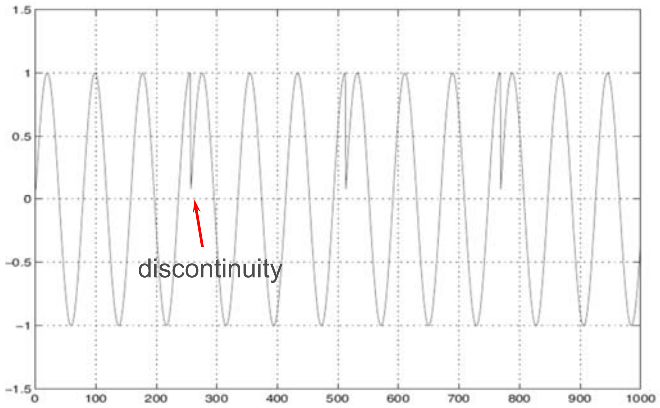

Spectral leakage

If the number of periods of the sampled signal periods is NOT an integer, as in the example below → the Fourier coefficients are non-zero at all harmonics (bins)!

This is referred to as *spectral leakage* → it looks as if the signal has long skirts → this is not to be confused with a proper noise floor!!

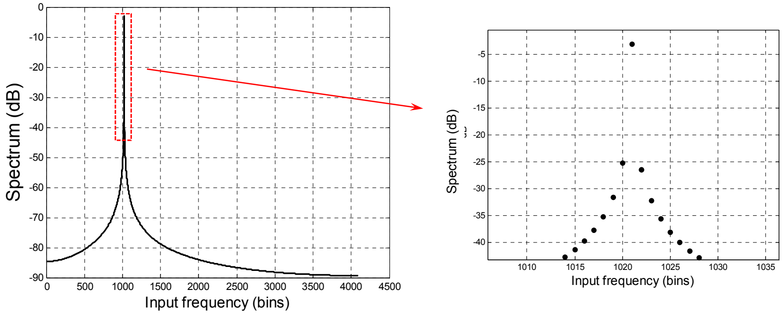

Windowing

Windowing addresses this problem by tapering the endings of the series:

$$x_w(k) = x(k)\cdot W(k)$$

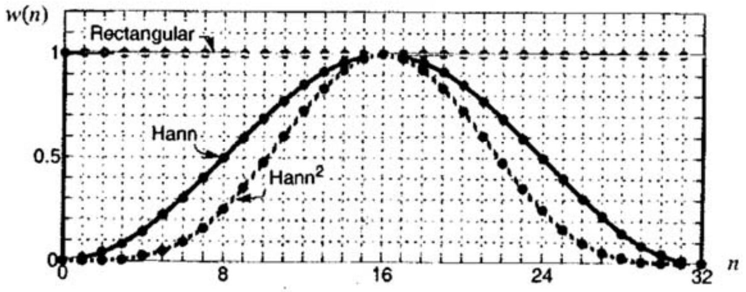

Useful tips

Nyquist-rate A/D simulations: make sure that the sequence of the samples is N-periodic (*coherent* sampling – more about it in a few slides); otherwise, use windowing

ΔΣ simulations: use always windowing!

With sine-wave inputs, avoid repetitive patterns in the sequence: the ratio between the sine-wave period and the sampling period should be an integer prime number – otherwise, noise is not spread over all bins, but rather concentrated over only a fraction of the bins → this applies in particular to Nyquist-rate converters, where, unlike in ΔΣ modulators, the same input results in the same q-error (no memory in the system)

Windowing example – rect/Hann/Hann^2 windows

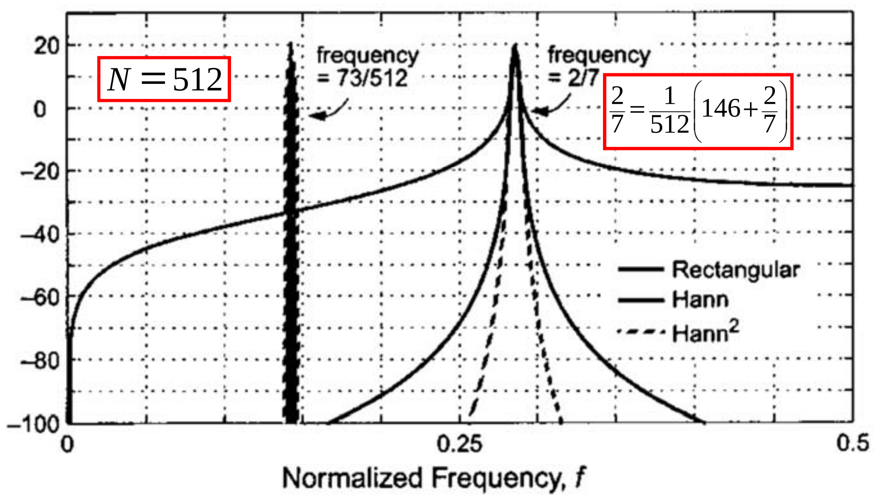

Hann window – signal bins

Rectangular window + integer number of signal periods → the signal is found over a single frequency bin (as expected)

Hann → the signal is spread over three bins → signal + two sidebands Further, the amplitude of the central bin is 6dB lower than with rectangular windowing

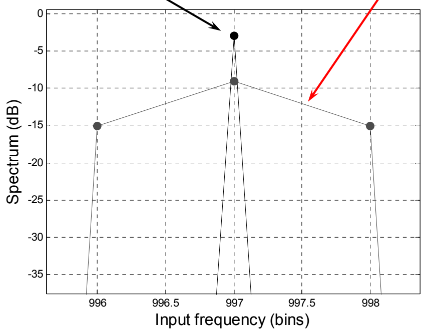

Is windowing avoidable?

Windowing is an amplitude modulation of the input, and as such introduces some undesired effects → e.g., a spike at the beginning/end of the sequence is completely masked

We have seen that windowing gives rise to spectral leakage of its own (i.e., even if the input is a sine wave, the spectrum is not a pure tone)

For input sine waves → to repeat: use coherent sampling, for which an integer number *k* of signal periods fits into the sampling window (moreover, in Nyquist-rate converters *k* should be chosen with care, if we want the q-noise to be truly white-like)

In this way, windowing can be avoided in Nyquist-rate converters (or more exactly, when q-noise is not shaped)

If q-noise is shaped → use windowing! (Hann is enough – more about this in a few slides)

Coherent sampling in detail

sampling period *signal period*

Coherent sampling: $m \cdot T_s = k \cdot T_{in}$

If: $m = a \cdot b$ and $k = a \cdot c$ → $a \cdot (b \cdot T_s) = a \cdot (c \cdot T_{in})$

i.e., in Nyquist-rate (non-feedback) converters the q-noise sequence repeats itself a times, identically → q-noise is not white

To avoid this, *m* and *k* should be prime to each other; in particular, if *m*=2^N, *k* should be odd:

odd number

sine-wave freq. $f_{in} = \frac{k}{2^N} f_s$ sampling freq.

# of samples

k and q-noise (Nyquist-rate case)

If k in $f_{in} = f_s \cdot k/2^N$ is not odd, the quantization noise repeats itself in exactly the same way two or more times across the sampling sequence → *q*-noise is not white, but shows repetitive patterns → q-noise is concentrated at certain frequencies!
The two simulations below use the same amount of samples, but notice how in the plot to the right the q-noise is much less white, with much higher peaks (possibly hiding signal harmonics from distortion)

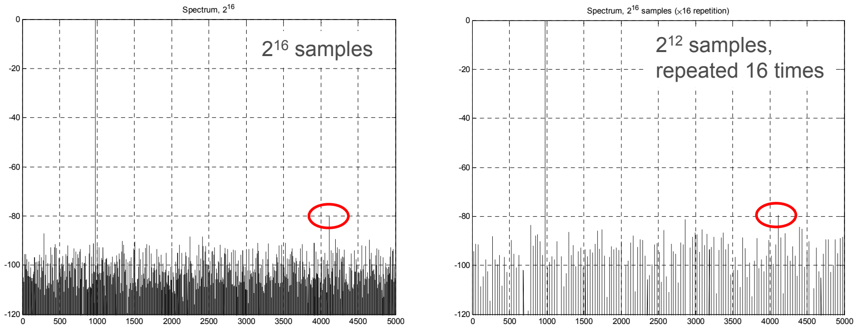

Noise and windowing in ΔΣ

The implicit rectangular window that we use when we select a finite sequence of samples from an infinite sequence is in fact a highly explicit window!

Multiplication in time is convolution in frequency → it is important to look at the frequency spectrum of the various windows! → rectangular window is discontinuous at endpoint → we expect much higher frequency content than Hann and Hann^2

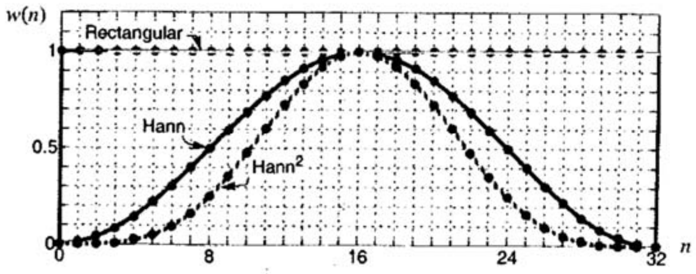

Noise and windowing in ΔΣ

Rectangular → *sinc* frequency response → high-frequency lobes approach a constant value (proportional to 1/N)

Hann → falls with 60dB/decade

Hann^2 → falls with 100dB/decade

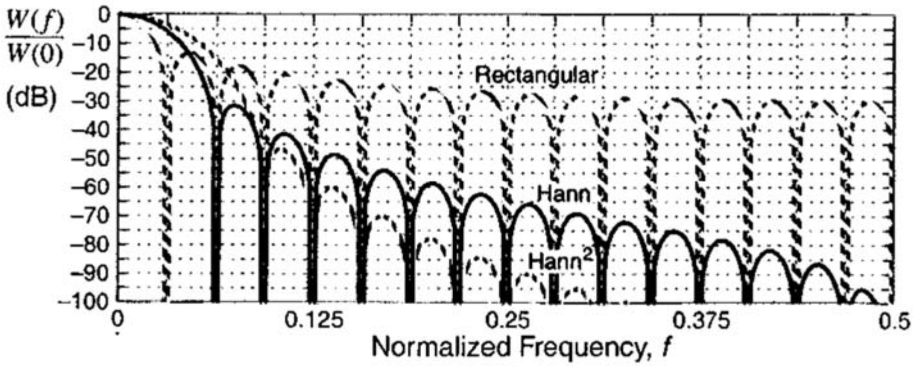

Noise and windowing in ΔΣ

The high-frequency spectrum of the rectangular window convolves with the high-frequency shaped noise of the ΔΣ modulator, folding back into the baseband and filling in the noise notch → dramatic noise deterioration!

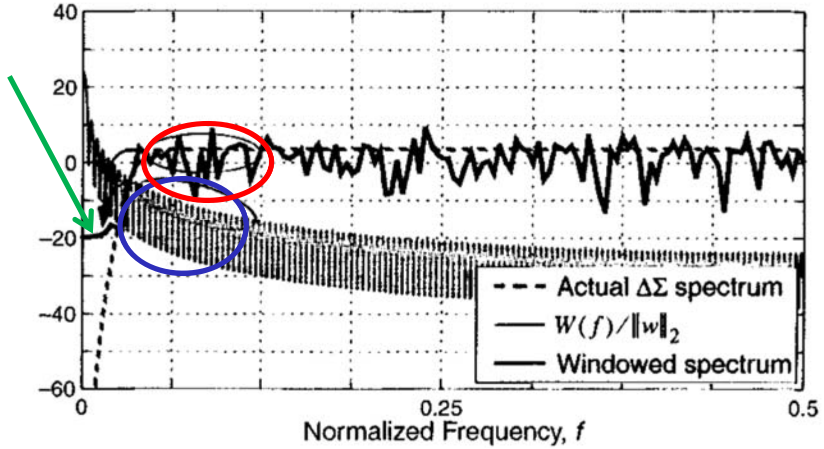

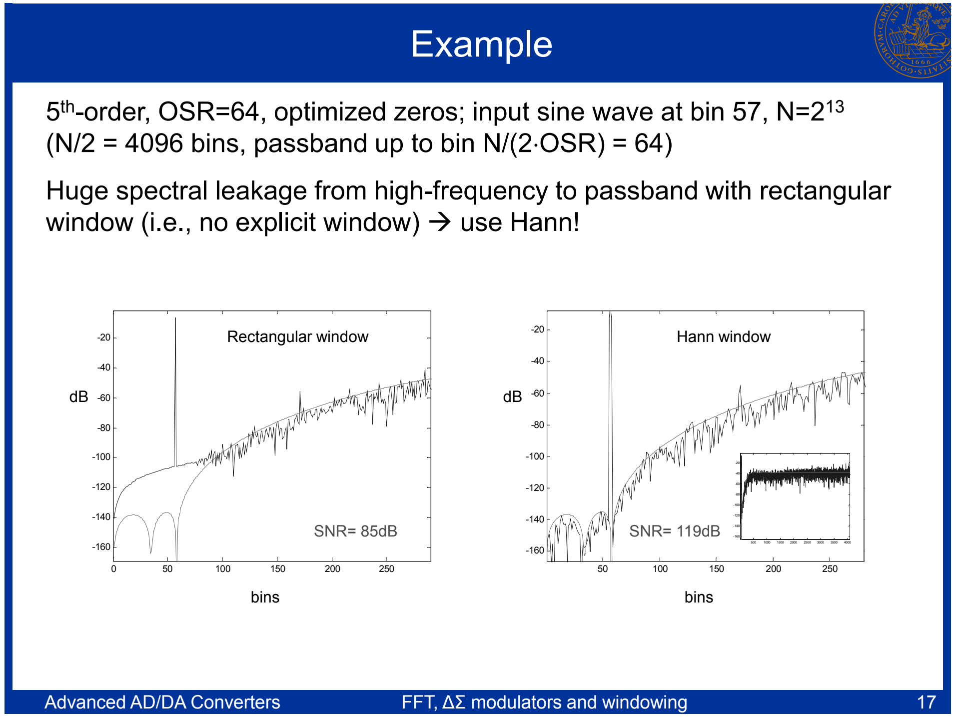

Example

5th-order, OSR=64, optimized zeros; input sine wave at bin 57, $N=2^{13}$ (N/2 = 4096 bins, passband up to bin N/(2·OSR) = 64)

Huge spectral leakage from high-frequency to passband with rectangular window (i.e., no explicit window) → use Hann!

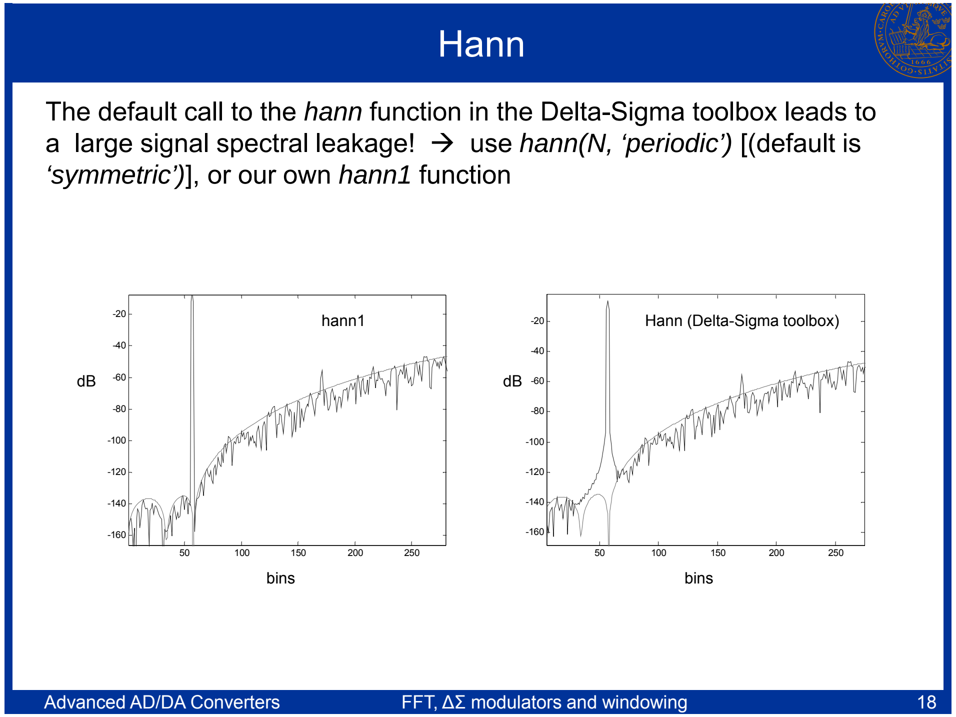

Hann

The default call to the *hann* function in the Delta-Sigma toolbox leads to a large signal spectral leakage! → use *hann(N, ‘periodic’)* [(default is *‘symmetric’*)], or our own *hann1* function

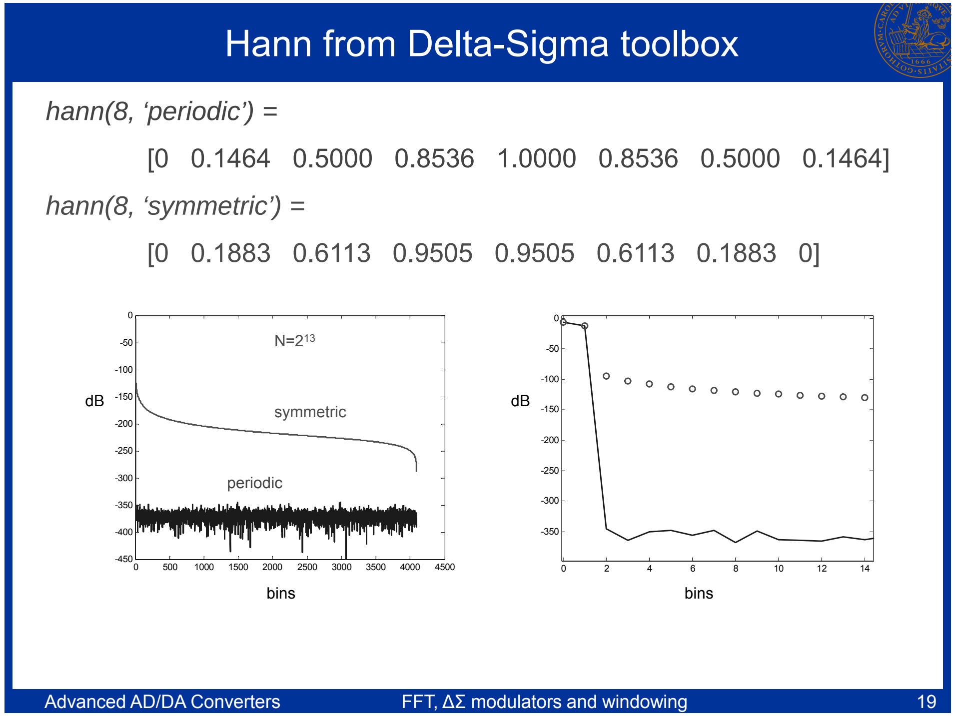

Hann from Delta-Sigma toolbox

hann(8, ‘periodic’) =

[0 0.1464 0.5000 0.8536 1.0000 0.8536 0.5000 0.1464]

hann(8, ‘symmetric’) =

[0 0.1883 0.6113 0.9505 0.9505 0.6113 0.1883 0]

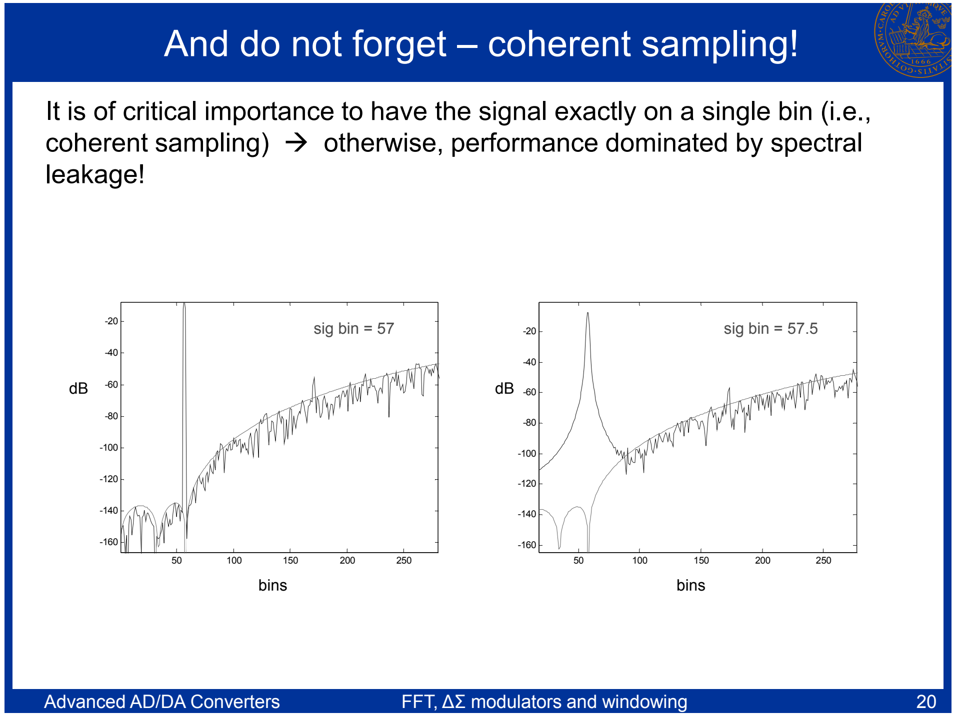

And do not forget – coherent sampling!

It is of critical importance to have the signal exactly on a single bin (i.e., coherent sampling) → otherwise, performance dominated by spectral leakage!

More on the FFT

The signal is usually placed in-band, and therefore occupies some of the in-band bins (typically, 3 with Hann) → there should be enough in-band bins that the error on the in-band noise is small → if the signal occupies 20% of the in-band bins, the error is below 1dB if the noise floor is flat → with 3 signal bins, we need 15 in-band bins → in all, we need N > 15·(2·OSR) = 30·OSR samples

The problem, however, is that the in-band noise is not flat, and it makes a difference where the signal is placed in the band itself!

Numerical experiments show that the SNR has a standard deviation of 1dB if N=64·OSR is used → recommended number of samples

N=64·OSR → an SFDR which is 10dB above the SNR can be detected – if not enough, a doubling of N increases the resolution by 3dB, thanks to the processing gain

---

More on the FFT

The height of the sine wave after the FFT is of course dependent on the kind of window we have used – for our windows we have

$$\max\left|W(f)\right| = W(0)$$

which means (since the output spectrum is given by the convolution of the non-windowed spectrum with $W(f)$) that the height of the peak of a sinusoid of amplitude A is A·W(0)/2 (since half of the energy goes into the mirror signal beyond Nyquist)

Spectrum plots: usually the FFT is scaled such that a full-scale input sine wave results in a 0dB peak (0dBFS) → if full-scale range is FS, the full-scale amplitude of a sine wave is FS/2 (since the whole peak-to-peak sine wave must be contained within FS) → the scaled expression of the PSD is therefore

$$S_x(f) = \left|\frac{\sum_{n=0}^{N-1} w(n)\cdot x(n)\cdot e^{-j2\pi fn}}{W(0)/2\cdot FS/2}\right|^2$$

More on the FFT

Alternative way of scaling the PSD, yielding a calibrated noise PSD (i.e. the noise floor is independent of window type and length):

$$S_x(f) = \frac{\left|\sum_{n=0}^{N-1} w(n)\cdot x(n)\cdot e^{-j2\pi fn}\right|^2}{\|w\|_2^2}$$

with $\|w\|_2^2 = \sum_{n=0}^{N-1} \left|w(n)\right|^2$ being the energy of the window – however, in this way the height of the signal (which is a power concentrated over one or few bin, and not a power spectral density spread out over all bins) depends on window type and length

Best option → use the first definition of $S_x(f)$, and give the associated *noise bandwidth (NBW)* of the window used (in fact, we use the same concept in measurements with a spectrum analyzer) → in this way, we can use the sine-wave-scaled FFT, and use the NBW to calculate the correct value of the noise power

---

Properties of Rect/Hann/Hann^2 windows

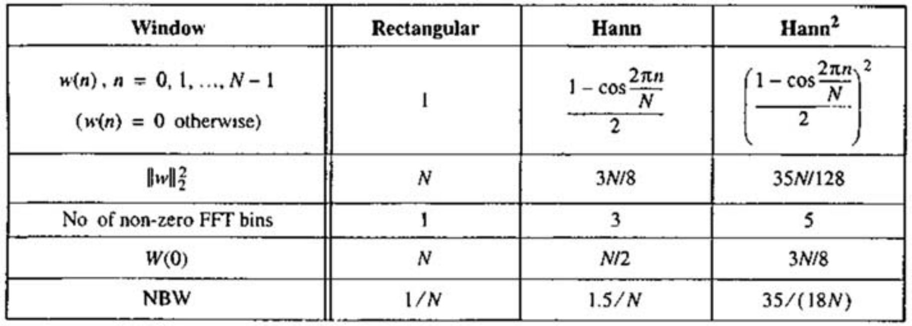

Window	Rectangular	Hann	Hann^2
$w(n)$, $n = 0, 1, \ldots, N-1$ ($w(n) = 0$ otherwise)	1	$\frac{1-\cos\frac{2\pi n}{N}}{2}$	$\left(\frac{1-\cos\frac{2\pi n}{N}}{2}\right)^2$
$\|w\|_2^2$	N	$3N/8$	$35N/128$
No of non-zero FFT bins	1	3	5
$W(0)$	N	$N/2$	$3N/8$
NBW	$1/N$	$1.5/N$	$35/(18N)$

$$NBW = \frac{\|w\|_2^2}{\left|W(0)\right|^2} = \frac{\|w\|_2^2}{\|w\|_1^2}$$

(assuming $w(n) \geq 0$, we have $\left|W(0)\right| = \|w\|_1$)

Example

Below: example with rectangular window (i.e. NBW=1/N), with both sine wave and noise having a power of 0dBFS
Top curve: noise floor of -15dBFS, which is the amount of power over a bandwidth NBW → the total noise power between DC and Nyquist (=0.5) is therefore the noise floor multiplied by 0.5/NBW (which, in the case of the rectangular window, is the exact inverse of the processing gain of the FFT)