

## FFT window functions


### Limits on FFT analysis

When using FFT analysis to study the frequency spectrum of signals, there are limits on resolution between different frequencies, and on detectability of a small signal in the presence of a large one.

There are two basic problems: the fact that we can only measure the signal for a limited time; and the fact that the FFT only calculates results for certain discrete frequency values (the 'FFT bins'). The limit on measurement time is fundamental to any frequency analysis technique: the frequency sampling is peculiar to numerical methods like the FFT.

### Limited measurement time

The first problem arises because the signal can only be measured for a limited time. Nothing can be known about the signal's behaviour outside the measured interval. We have to assume something about the signal outside the measured interval, and the Fourier Transform makes an implicit assumption that the signal is repetitive: that is, the signal within the measured time repeats for all time.




Most real signals will have discontinuities at the ends of the measured time, and when the FFT assumes the signal repeats it will assume discontinuities that are not really there.

Since sharp discontinuities have broad frequency spectra, these will cause the signal's frequency spectrum to be spread out.

### Spectral leakage

It is easy to gain an insight by thinking about the special case of a pure sine wave. This has a frequency spectrum which is a single spectral line: but the frequency spectrum calculated by the FFT will show a spread out line. Each spectral line will be spread out in the same way.



The spreading means that signal energy which should be concentrated only at one frequency instead leaks into all the other frequencies. This spreading of energy is called 'spectral leakage'.

Since spectral leakage is related to discontinuities at the ends of the measurement time, it will be worse for signals that happen to fall such that there are large discontinuities.

Some signals may, by coincidence or by design, fall in such a way that there happens to be no discontinuity at the ends of the measurement time: for these signals the effect of spectral leakage may be lessened.

For example a pure sine wave sampled for an exact number of cycles would match up quite correctly when made repetitive: the repetitive signal would be exactly the same as the 'real' signal and so no spectral leakage would occur.

Another example would be a signal which fell smoothly to zero at each end of the measurement interval: such a signal would have no discontinuities when made repetitive, and so would not suffer from spectral leakage.

Such special cases are infrequent,

but they can be arranged: for example frequency analyses are often made by tuning a stimulus frequency precisely so that its frequency exactly fits the measurement interval.

Spectral leakage is not an artifact of the FFT, but is due to the fact that the signal was measured only for a finite time. For a sine wave to have a single line spectrum it must exist for all time. Any practical method of calculating the frequency spectrum of a signal suffers from spectral leakage due to the finite measurement time.

Spectral leakage is not related in any way to the fact of having sampled the signal, but only to the finite measurement time.

Spectral leakage causes at least two distinct problems.

First, any given spectral component will contain not just the signal energy, but also noise from the whole of the rest of the spectrum. This will degrade the signal to noise ratio.

Second, the spectral leakage from a large signal component may be severe enough to mask other smaller signals at different frequencies.


### Windowing

In effect, the process of measuring a signal for a finite time is equivalent to multiplying the signal by a rectangular function of unit amplitude: the rectangular function lasting for the duration of the measurement time.

The signal is measured during a finite measurement time or 'window'. This idea leads to the rectangular function being called a 'rectangular window'.

The effects of spectral leakage can be reduced by reducing the discontinuities at the ends of the signal measurement time.

This leads to the idea of multiplying the signal within the measurement time by some function that smoothly reduces the signal to zero at the end points: hence avoiding discontinuities altogether.



The process of multiplying the signal data by a function that smoothly approaches zero at both ends, is called 'windowing': and the multiplying function is called a 'window' function.

It is easy to analyse the effect of a window function: the frequency spectrum of the signal is convolved with the frequency spectrum of the window function.


### Spectral leakage

One way to visualise spectral leakage is as spreading of the frequency components.

Each frequency should contribute only to one FFT bin but spectral leakage causes the energy to be spread by the window function so it contributes to all other FFT bins.

The contribution is weighted by the window function centred at the frequency component and evaluated at the FFT bin.


In the special case of the rectangular window (that is, no window at all), the window function is 1 in the interval and 0 outside the interval. Its Fourier transform is known as the 'sinc' function, or more formally the 'Dirichlet kernel'.



The shape of the Fourier transform of a window function is sometimes called the 'kernel'.

(Confusingly, the Fourier Transform of a window function is also often called the 'window function': we have to judge from context whether we are talking about the frequency or the time domain function.)

### The FFT as a series of filters



Another productive way to visualise spectral leakage is to regard the FFT as equivalent to a series of filters, centred on each spectral sample.

The filter's frequency response is the shape of the window function.

This is the inverse of the spectral leakage model. Each FFT bin includes contributions from all other frequencies in the bandwidth of the filter, weighted by the window function. This will include contributions from broad band noise as well as narrow band signals at other frequencies.

### Detection or resolution?


There are two common situations: detection of a spectral component in the presence of broadband noise; or distinguishing between narrow band spectral components. The choice of window function may be different in the two cases of detection or

resolution.

To help in choosing a suitable window function some quantitative measures are needed.

### Equivalent noise bandwidth

A given FFT bin includes contributions from other frequencies including accumulated broadband noise. To detect a narrow band signal in the presence of noise, we want to minimise the noise. This can be done by using a narrow bandwidth window function.



The Equivalent Noise Bandwidth (ENB) of the window measures the noise performance of the window. It is the width of a rectangle filter which would accumulate the same noise power with the same peak power gain. This is a fruitful concept, and easy to visualise.

### Coherent power gain

'Coherent Power Gain' measures the reduction in signal power due to the window function suppressing a coherent signal at the ends of the measurement interval.

Again, this measurement relates to the model of the FFT as a bank of filters matched to each frequency sample.

For an ideal discrete line frequency component, the 'noiseless' signal contribution to the FFT bin is proportional to the signal amplitude. The proportionality factor is the sum of the window terms, which is just the DC gain of the window. The 'Coherent Gain' is the square of this, or in other words the

coherent power gain is the square of the sum of the window terms.

For a rectangular window the DC gain is  $N$ , the number of terms in the window: but for any other window the DC gain will be reduced because the window goes smoothly to zero at the ends of the measurement time.

The reduction in DC gain is important because it represents a definite scaling of the amplitudes of the frequency spectrum which requires correction for any absolute measurements to be correct.

Coherent gain is usually normalised by dividing by  $N$ , so that the normalised coherent gain of a rectangular window is 1.

### Processing Loss

'Processing Loss' is the ratio of input signal to noise to output signal to noise, which is the Coherent Power Gain divided by the noise power.

For a signal made up of an ideal discrete line frequency component polluted by white noise, the Processing Loss is, not unexpectedly, the reciprocal of the Equivalent Noise Bandwidth (ENB).

### Scalloping loss

The analysis of Equivalent Noise Bandwidth and Processing Loss so far related only to signal components whose frequencies exactly match the FFT bins.

The model of the FFT as a series of filters centred on the FFT bins suggests that frequencies that do not coincide with the FFT bins will be attenuated as the filter's response falls off away from the centre frequency.

This effect will vary as the signal frequency ranges between the adjacent FFT bins, and is called 'the picket fence effect' or 'scalloping'.

It is a reasonable assumption that the worst case occurs when the

signal frequency lies exactly half way between FFT bins.

'Scalloping Loss' is defined as the ratio of coherent gain for a signal frequency component located half way between FFT bins, to the coherent gain for a signal frequency component located exactly at an FFT bin.

This is just the ratio of the window function's value one half a frequency sample off centre, to its value at the centre frequency.


### Worst case processing loss

Worst case processing loss is defined as the sum of Scalloping Loss and Processing Loss.

This is a measure of the reduction of output signal to noise ratio resulting from the combination of the window function and the worst case frequency location. It is of course related to the minimum tone that could be detected in broadband noise.

### Spectral leakage

The calculated spectrum is not only affected by broadband noise, but also by the narrow band frequency components either of noise or of the signal.



A frequency component at frequency  $\omega_0$  will contribute to the calculation at another frequency  $\omega_a$  proportionally to the gain of the window's Fourier transform, centred at  $\omega_0$  and measured at  $\omega_a$ . This is usually referred to as 'spectral leakage'.

Spectral leakage can change the calculated amplitude and position of

a spectral estimate.

The effects of spectral leakage are worst when detecting small signals in the presence of nearby large signals.

To minimise the effects of spectral leakage, a window function's FFT should have low amplitude sidelobes away from the centre, and the fall off to the low sidelobes should be rapid.

The peak sidelobe level is a useful indicator of how well a window function suppresses spectral leakage: so is the rate of fall off to the sidelobes.

### Minimum resolution bandwidth

An interesting measure is the minimum separation needed between two frequency components of equal amplitude, so that they can be resolved.

By resolved, is meant that there should be a local minimum between the two peaks.

The 'rule of thumb' for resolvability is the width of the window at the half power points (the 3 dB bandwidth): because two frequency components of equal strength show a single peak if separated by less than their 3 dB bandwidth, and so cannot be separated.

But this assumes incoherent addition. There is an important problem with this criterion when applied to the FFT because the addition which is involved with the FFT is coherent, not incoherent

The FFT output is the coherent addition of frequency components, weighted by the window function at each frequency. Because of the coherence, the 6 dB bandwidth defines resolution rather than the 3 dB bandwidth.

If two frequency components are involved the sum of the window functions at the crossover point (halfway between the peaks) must

be smaller than the individual peaks if the two peaks are to be resolved. So the gain from each window function must be less than 0.5 (or 6 dB).

### Some window functions

The table lists, for various window functions, the following parameters:

<b>Sidelobe level</b>	The attenuation to the top of the highest side lobe
<b>Fall off</b>	The asymptotic rate of fall off to the side lobe
<b>Coherent gain</b>	The normalised DC gain
<b>Equivalent noise bandwidth</b>	The bandwidth of a rectangular filter which would let pass the same amount of broadband noise
<b>6 dB bandwidth</b>	The bandwidth in which the window function falls by 6 dB
<b>Worst case processing loss</b>	The ratio of input signal to noise to output signal to noise, including scalloping loss for the worst case frequency

Window function	Sidelobe level (dB)	Fall off (dB per octave)	Coherent gain	Equivalent noise bandwidth (bins)	6 dB bandwidth (bins)	Worst case processing loss (dB)
Rectangular	-13	-6	1.00	1.00	1.21	3.92
Triangular	-27	-12	0.50	1.33	1.78	3.07
Hanning	-32	-18	0.50	1.50	2.00	3.18
Hamming	-43	-6	0.54	1.36	1.81	3.10
Poisson (3.0)	-24	-6	0.32	1.65	2.08	3.64
Poisson (4.0)	-31	-6	0.25	2.08	2.58	4.21
Cauchy (4.0)	-35	-6	0.33	1.76	2.20	3.83
Cauchy (5.0)	-30	-6	0.28	2.06	2.53	4.28
Gaussian (3.0)	-55	-6	0.43	1.64	2.18	3.40
Kaiser-Bessel (3.0)	-69	-6	0.40	1.80	2.39	3.56
Kaiser-Bessel (3.5)	-82	-6	0.37	1.93	2.57	3.74