

FFT Properties 3.5

Frequency analysis tutorial

Written by: Janez Atmapuri Makovsek
Dew Research and Development
Date: 10th September 2005
Revision: 1.3

1	INTRODUCTION.....	2
2	SIGNAL TYPE.....	3
3	THE PERIODICITY	8
4	ALIASING OR AMBIGUITY	9
5	WINDOWS.....	10
5.1	HANNING WINDOW	11
5.2	FLAT TOP WINDOW	14
5.3	BLACKMAN WINDOW.....	16
5.4	HAMMING WINDOW	17
5.5	EXPONENTIAL WINDOW	18
5.6	SUMMARY ON WINDOWS	18
6	DC COMPONENT	19
7	BEATING	21
8	AMPLITUDE MODULATION AND SUM AND DIFFERENCE FREQUENCIES.....	23

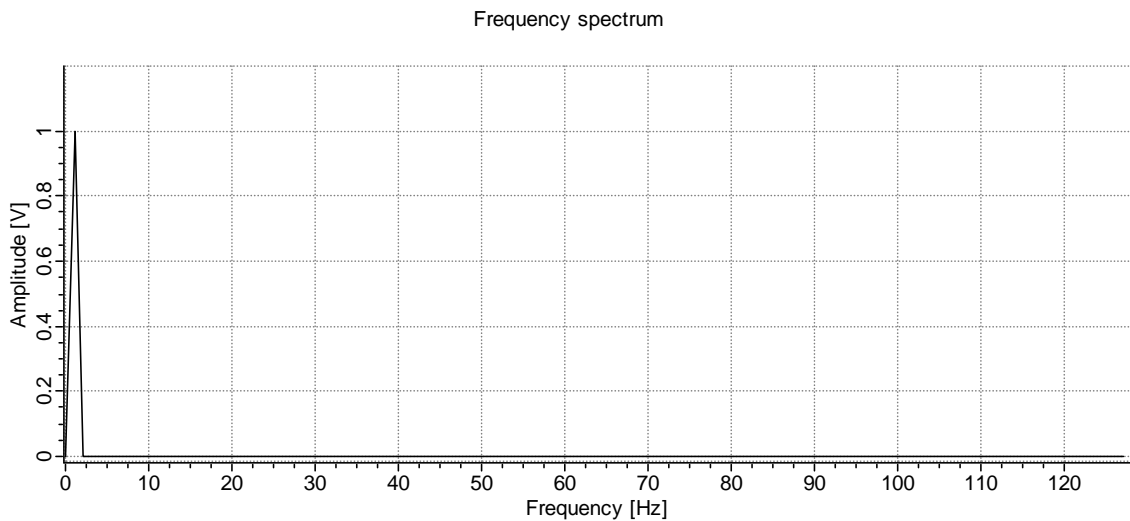
1 Introduction

The tutorial is based on the experimental point of view. It is meant to be used together with the software: [FFT Properties 3.5](#). This tutorial does not include any formalized definitions and formulas, because those are available in many DSP books. The tutorial focuses on application of the FFT based frequency analysis and related issues. Any special pre-knowledge is not required, but only the synthesis of theory and application will give the reader a chance to make the best of it. The program uses a number of known signals to illustrate what can go wrong with discrete analysis of continuous signals. My strategy is to show you examples of what can go wrong in order to elucidate the theory and methods which allow accurate signal transformation.

2 Signal type

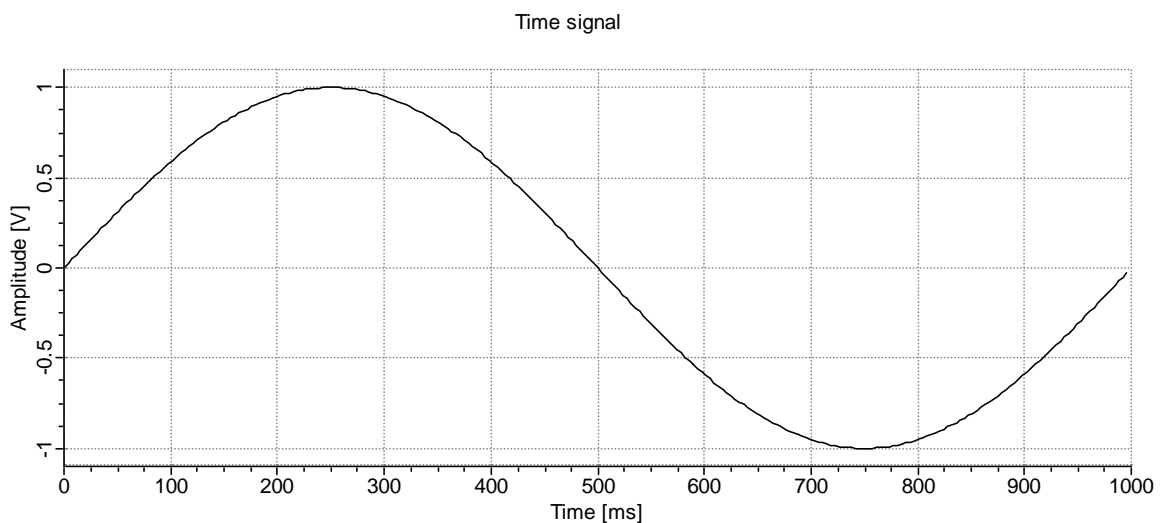
Most real world signals are continuous in nature. This is even true of very short pulses lasting only a tiny fraction of a second. The computer or, more accurately the soundcard and its analog to digital converter (ADC), digitizes these signals by very quickly sampling them as they occur. How quickly is known as the Sampling Rate which is measured in terms of samples per second. So a sampling rate of 44100, which is probably the highest rate of your soundcard's ADC produces 44100 different snapshots of the signal each second.

When the FFT Properties starts, you should switch to the FFT Panel, check the Sound check button to Off and set Sample time to 1000 ms, then Switch to the Signal Generator panel and set the amplitude and the frequency of the signal one to 1. You should now be looking at this (Picture 1):



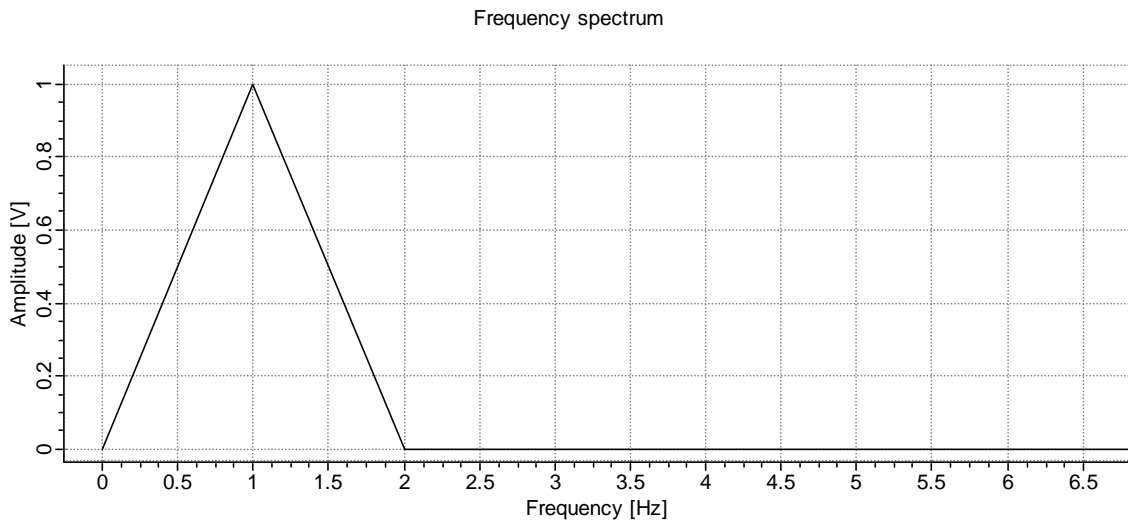
Picture 1 Frequency spectrum of a tone with 1 a frequency of 1Hz.

At first it may seem as though there is not much of significance about the upper chart. It is a frequency spectrum of the signal below (a time signal) which has a frequency of 1Hz and an amplitude of 1V. The graph of this time signal below (Picture 2) that shows just one cycle is also very straightforward.



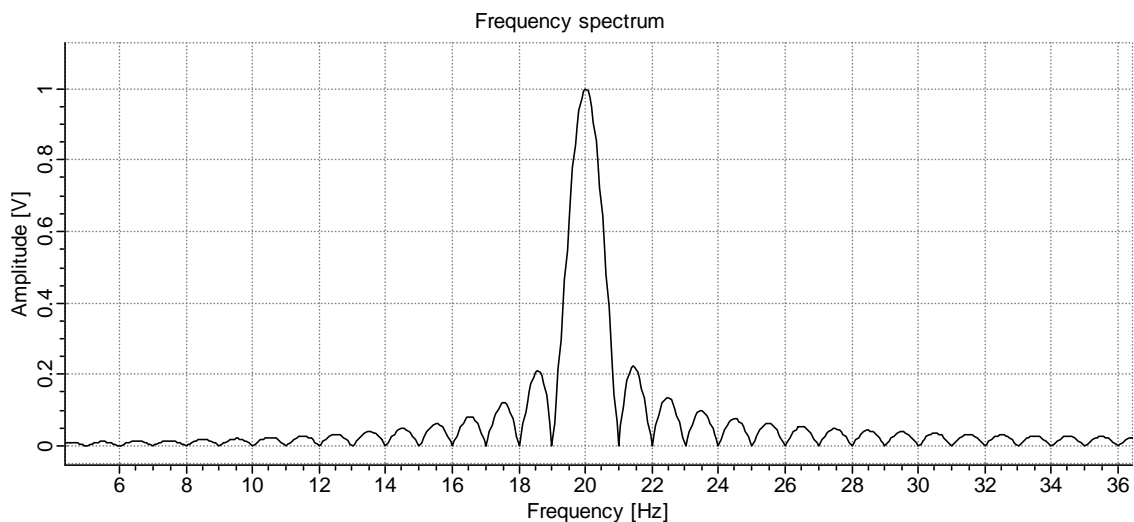
Picture 2 Time signal of a 1Hz tone signal.

However, closer observation (by using zoom) shows that the frequency spectrum seems to have frequencies that are not actually present in the original signal¹.



Picture 3 Zoomed-in frequency spectrum.

It looks as if there are some frequencies present in the original signal between 0 and 1 Hz and 1 and 2Hz. This is clearly some artifact of the process of producing the spectrum having to do with the RESOLUTION at which the spectrum is sampled. The spectrum resolution is displayed over the spectrum and yields: 1Hz/lin. That means that there is one point of data for each integer frequency point from left to the right. Therefore we have points calculated at: 0, 1, 2 , 127. This is the cause for the triangle and apparent additional frequencies. The mystery is what amplitudes are between the currently calculated lines. There is a way to sample the spectrum at a desired number of lines and, more importantly using many more lines. That parameter is defined on the FFT Control page and is called zero padding. Increase the frequency to 20 Hz and set zero-padder to 16 and observe the result (zoom-in around 20 Hz, Picture 4).



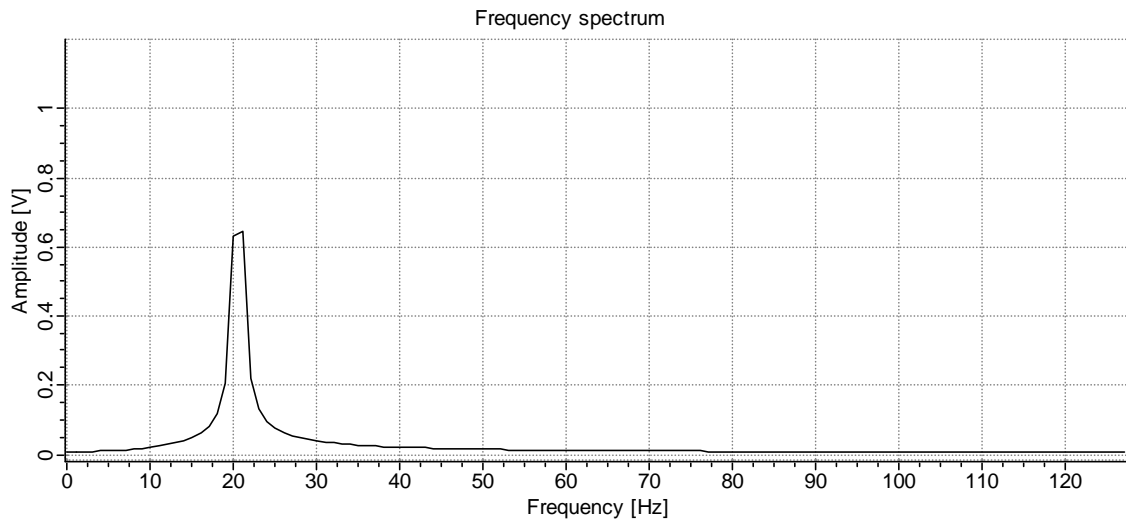
Picture 4 Leakage in the frequency spectrum

We now see what was between the lines for this 20 Hz. signal. In the original setting (zero padder = 1) we saw a spectrum sampled only at each integer frequency. We can see from the upper picture that only at each integer frequency is the amplitude zero. We were misled by the original spectrum. The problem is obvious. We have a whole lot of frequencies displayed that are not present in the signal at all. An important part of FFT analysis deals with methods on how to reduce these artifacts called

¹ You can zoom-in, by clicking and dragging a rectangle down and right. You can zoom out by dragging the rectangle up and left. You can pan the picture when zoomed in, by dragging it.

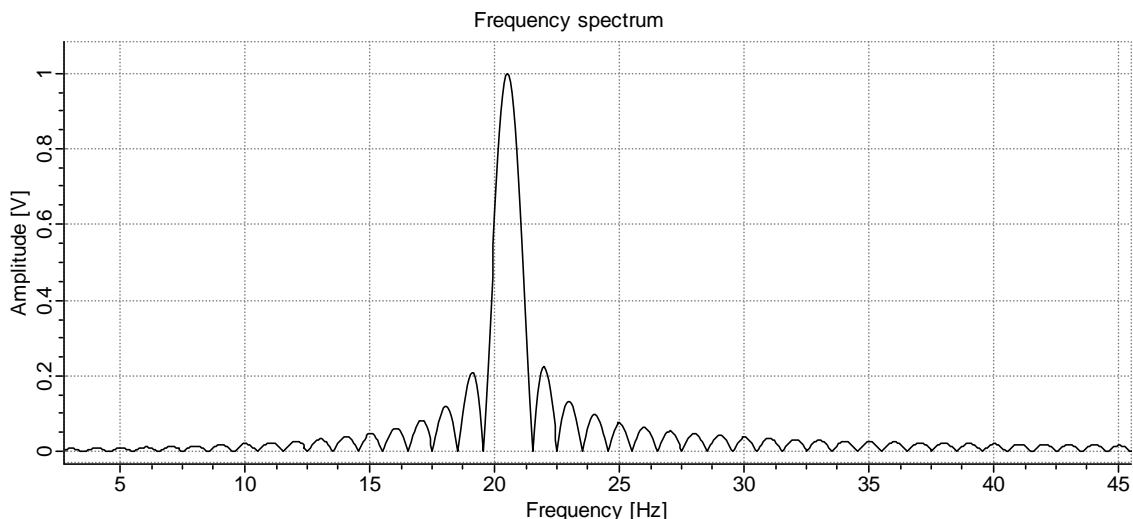
leakage which occurs because we are using discrete values to represent a continuously changing signal. This artifact called LEAKAGE is explained in the “The theory and praxis of DFT” (Another document included with FFT Prop) and is also referred to as Gibbs phenomena.

It is always good to have zero padding set to at least 2. This ensures that we will obtain a more accurate estimate of the frequency and of the amplitude of the signal. Now let's change back to zero padding = 1. What would happen if the actual frequency were not to coincide with the lines at which we have calculated the amplitude? Would we still be able to see the amplitude of that signal? Let us change the original signal frequency to 20.5 Hz by entering the value directly in the edit box.



Picture 5 Frequency spectrum of a tone with frequency 20.5Hz

It seems that there is a whole bundle of new frequencies (Picture 5). Besides, the amplitude does not show the actual amplitude of the signal. It has been greatly reduced. Be warned. This is also called leakage, but is of somewhat different from that previously illustrated. The energy is still there where it was. This is all the consequence of the sampling of the spectrum: how closely together do we calculate lines. Again, increase zero padding to 16. (Picture 6)



Picture 6 Zero padded frequency spectrum of frequency 20.5 Hz.

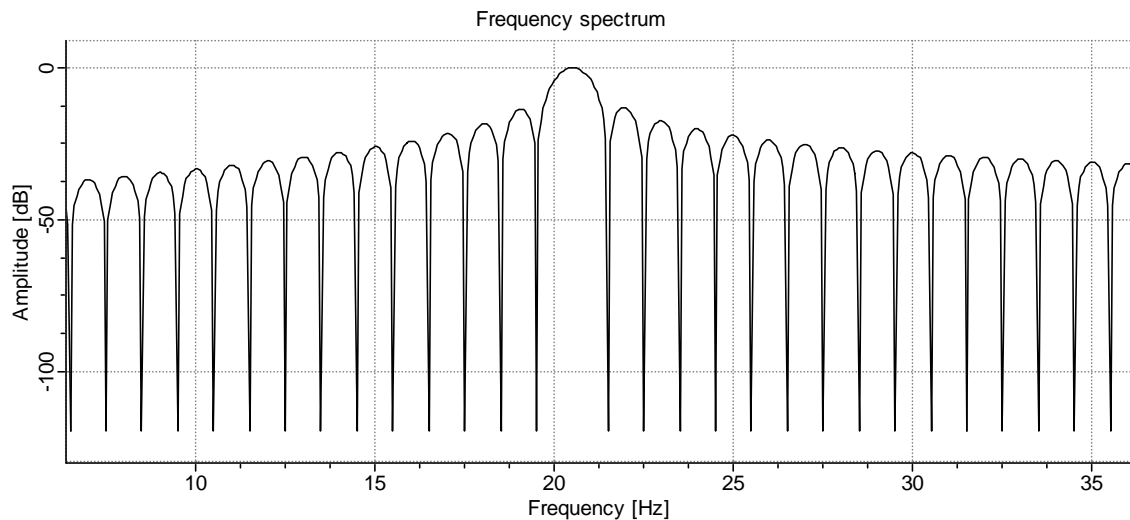
We get almost the same picture as with 20Hz, except that the peak is now at 20.5 where we set it. There are two good things to be found here:

1. The amplitude is now accurate

2. The frequency is now accurate

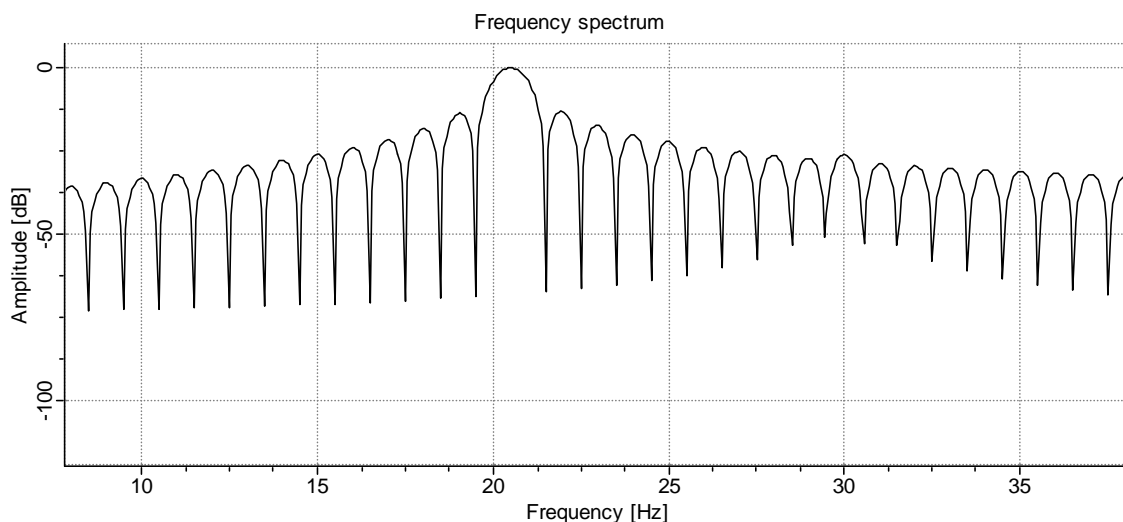
Why did we get so distorted a picture before? We were sampling the spectrum at many fewer points and this produced a misleading picture.

Let's take a look at the same picture with a logarithmic scale (Picture 7).² The whole spectrum is polluted across entire bandwidth with frequencies that are not present in the original signal at all.



Picture 7 Logarithmic frequency spectrum

That is really something to worry about. What if we had another frequency present in the signal, let's say at 30 Hz but with amplitude of 100 times smaller? Set Signal1 amplitude = 100 and Signal2 amplitude = 1 and the frequency of signal2 = 30 Hz. Here is the spectrum (Picture 8).



Picture 8 Logarithmic frequency spectrum of two frequencies.

As we expected, the original frequency is polluting the whole spectrum and the 100 times smaller frequency is barely noticeable. If we had more of such small signals, we wouldn't be able to really to detect them.

But you can also look at it the other way. If that distortion would not be present, then it would be necessary to calculate the amplitude at the known frequency. (For example 4.5Hz) In the real world

² You can switch to the logarithmic scale, by pressing the button captioned "Log" located top right just next to the frequency spectrum display. You can switch back to the linear scale by pressing the button caption "Lin". The logarithmic scale shows the data in decibels.

that might as well be : 4.56783231.. . It would be almost impossible to hit exactly! So in this way, we at least get the information that there is something. What actually is there is for us to find out. Usually we are interested in:

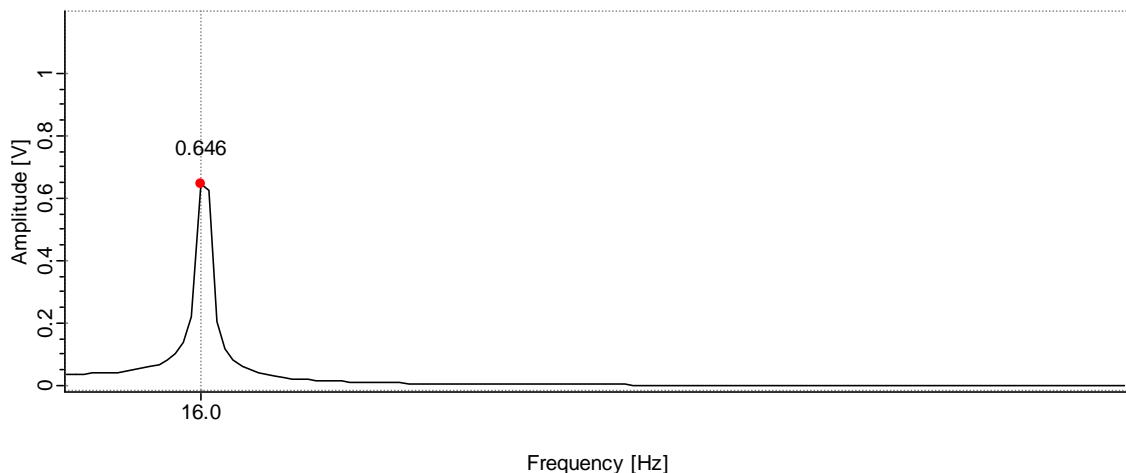
- What is the accurate amplitude of the given frequency?
- What is the frequency of the peak present in the spectrum?

Both of these things may be compromised, as we have seen. If the frequency of the signal falls between two lines, and we do not apply zero padding (spectrum oversampling) it is almost impossible to say what the actual frequency of that signal is or what the actual amplitude is. And besides, how do we know, that the frequencies are actually those that we can see? As far as we know, that could just as well be one misaligned (non-integer period) frequency instead of more aligned ones.

Switch again back to a non-zero padded signal by setting the zero padding to 1 and switch to linear scale (use lin button just above the log button next to the spectrum graph). Also make sure that the button labeled with "S" is pressed.³ Set Sample time on the FFT Control panel to 1050 ms. This is necessary so that we can increment the frequency in steps of 1 Hz and not get a an integer number of periods within the time signal window. Now place the cursor in the frequency edit box of Signal1 and use the Up/Down arrow keys on your keyboard to change the frequency.⁴

Despite the fact, that the amplitude of the original signal is constantly 1, that peak in the spectrum will occasionally fall by as much as 36% below 1. The graph below (Picture 9) shows a spectrum of signal with frequency of 16.5 Hz and Amplitude of 1V with the Sample time (window length) of 1000 milliseconds. The maximum of the spectrum was found at the 16 Hz. This is due to the fact that lines are calculated only at: 0, 1, 2, ..., 15, 16, 17..... [Hz] when zero padding is 1.

Frequency spectrum



Picture 9 Marked peak for a frequency of 16.5 Hz and amplitude 1.

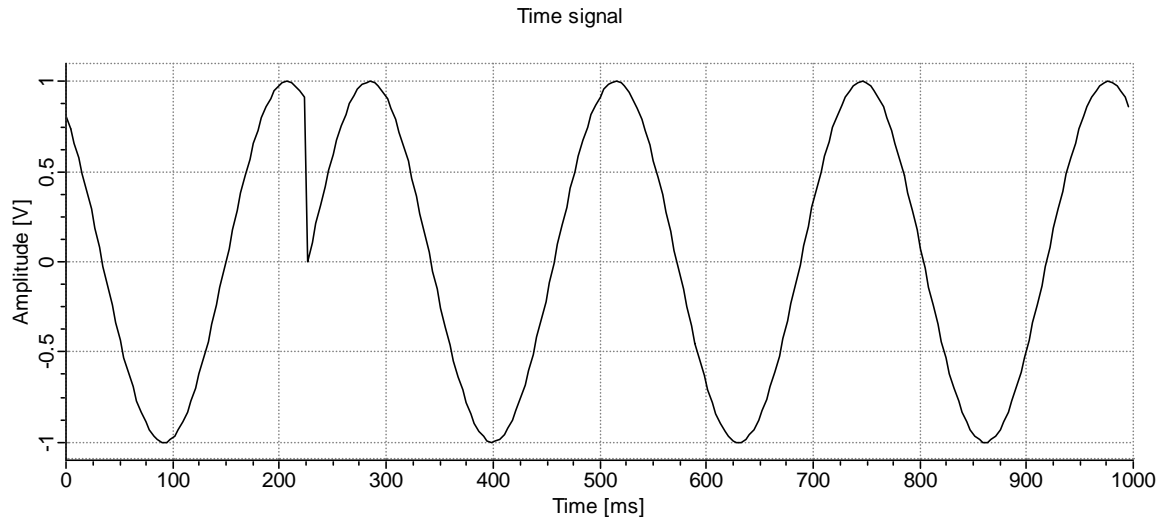
Different methods have been developed in the past to make such analysis easier, more accurate and faster. The most well known are Windows (not to be confused with the ubiquitous Microsoft OS or panes of glass : the analogy is that we need a means of seeing into something which doesn't distort that something...). All frequency analyzers offer at least three Window functions. But more on Windows and Windowing later on.

³ The "S" button enables auto-scaling with peak hold. The program only increases the scale, if the spectrum values exceed the display, but never decreases them. To reset the Y axis scale press the button again.

⁴ The frequency edit box can have its step defined with the Step drop down box located lower on the signal generator panel. When using the Up/Down arrow keys the initial frequency will be used only, if the value in the edit box is an integer.

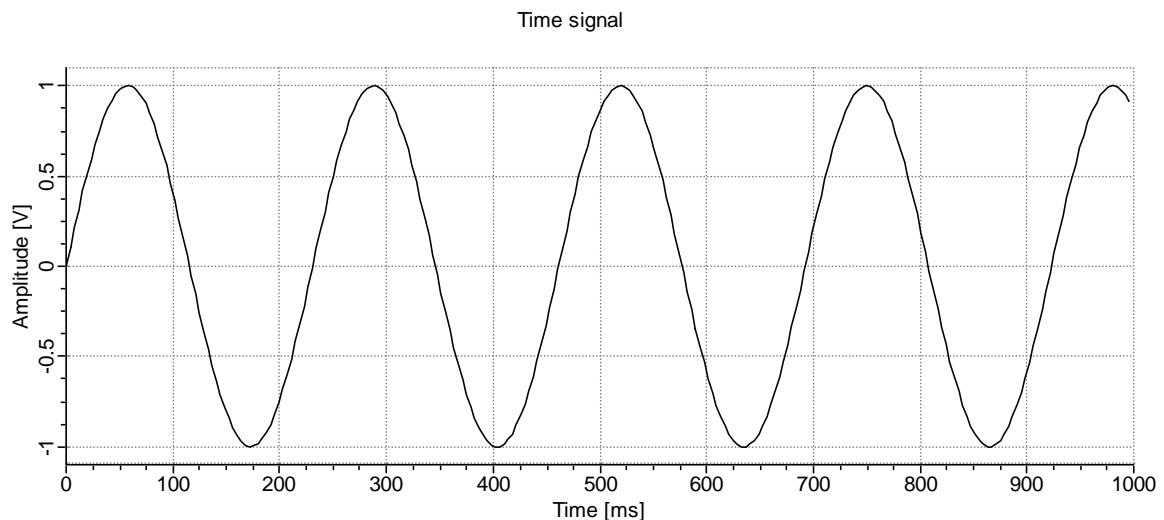
3 The periodicity

Although we have explored the effects of leakage, it is still not entirely clear where and why leakage originates. FFT assumes that the signal is periodic in time. What does that mean? If the number of periods (single cycles of the wave from zero to zero) is an integer number in the given time signal, then the next piece of the infinite (periodic) series would start exactly at the point where a preceding cycle finished. However, if the number of periods per time signal were not integer numbers, we would see breaks like this (Picture 10).



Picture 10 Non-integer number of periods per analysis window

We can see a discontinuity at approximately 220 ms. If we were to shift the time signal to the left (signal can be rotated with a 'Rotate' button on the Time Signal control panel), so that the anomaly reaches 0 ms, we would see this (Picture 11):



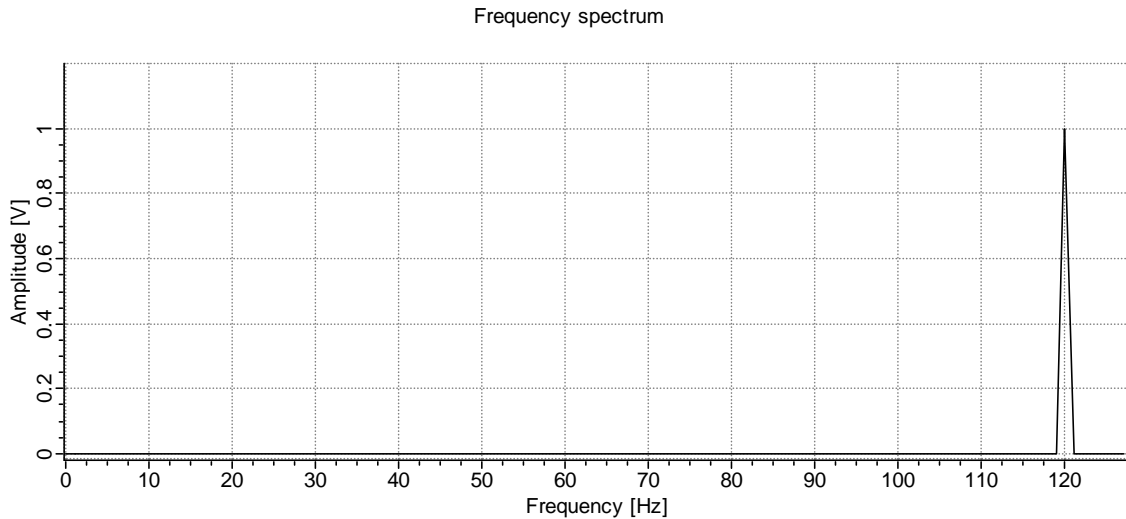
Picture 11 Non-integer number of periods per analysis window and rotated

Notice, how the amplitude at zero [ms] is 0 V and at 1000 [ms] the amplitude is approx. 0.9V. The first picture (Picture 10) was achieved by simply rotating the signal displayed on the second chart (Picture 11). The points from the end of the time series were inserted at the beginning and the whole reason for the leakage became obvious. The sharp step causes the additional frequencies. The first and the second chart have exactly the same frequency spectrum. This is because we are both times looking at the same signal, from the FFT point of view. If you set: Signal1 to 3.666 Hz (enter the value directly in the edit box in the Signal Generator) at 256 samples (FFT Control panel samples edit box) and

change the Rotate up-down button (located on the Time Signal panel) the time signal will be rotated by a given number of samples. The interesting thing to observe is the frequency spectrum, which does not change, despite the obvious visual changes in the time signal.

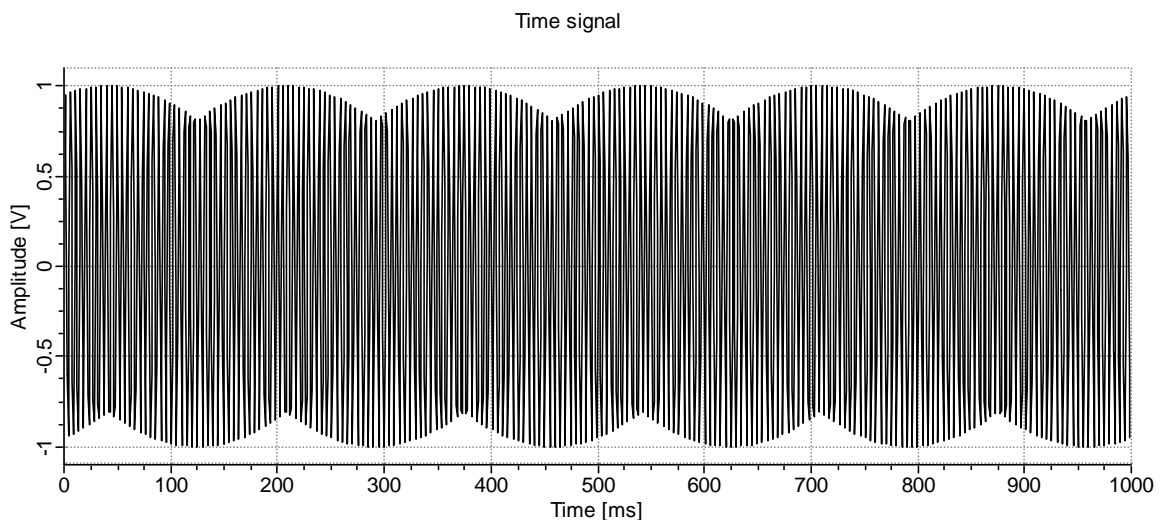
4 Aliasing or ambiguity

Aliasing is perhaps the most misleading effect in frequency analysis, since it can be the cause of the biggest misinterpretations. It is called "aliasing" because one thing is represented as something else in a way intended to mislead.



Picture 12 Frequency spectrum for frequencies 120Hz, 136Hz, 376Hz...

We seem to have one signal⁵ at a frequency of 120 Hz (Picture 12). But is that really so? Change the frequency of the signal 1 to 136 Hz. That peak should we would expect disappear from the scale. But it does not. When it reaches 127 Hz, the peak starts to travel back and when we reach the 136 Hz, the peak is again at 120 Hz. It gets much worse. Set the frequency of signal 1 = 376 Hz. Again the same spectrum. This effect repeats itself to infinity. The FFT algorithm does not cause this. To the contrary: it is a consequence of the way we sample discretely the original continuous signal.



Picture 13

⁵ Sample time edit box on the FFT Control panel should be set to 1000ms, Signal1 frequency to 120Hz and amplitude to 1.

The main problem that we have to deal with is: how to make sure that the frequencies in our spectrum are not from somewhere above 300kHz and are shown misleadingly as a frequency at 120Hz. The solution is very simple. All the frequencies above the upper limit of the spectrum must be attenuated in the signal "before" we digitize it. This procedure is called a low-pass filtering because it passes lower frequencies while excluding higher ones and must be applied to the signal before it enters the sound card.

Another related matter of interest is how aliasing is reflected in the time domain. Let's take a look at the Picture 13.

We have one very high frequency signal (1024 samples, 206Hz). But it seems that there are some low frequencies present. (Close to 1.5 Hz). This is not caused by the way the lines are charted. This is a direct consequence of the sampling. A very high frequency is sampled at points that have these values. You can even zoom in and the signal will retain its form. This example is not like the aliasing we have previously examined, except that it is a good example of how with FFT high frequencies may appear as low ones.

Many times, when analyzing real-world signals, we can see a lot of "noise", random like frequencies across the spectrum. Unfortunately, this type of noise could be an indication of improperly set low-pass filters. And what we see as noise, are actually very high frequencies wrapped around the upper and lower bounds of the amplitude spectrum. When sampling your data, you should make sure that the sampling frequency is higher than two times the highest frequency of interest that is physically present in the signal being sampled. This is called the Nyquist frequency. If you don't do this, then an anti-aliasing filter before the signal enters the system is mandatory. If such precautions are not taken then it is many times almost pointless looking at the data.

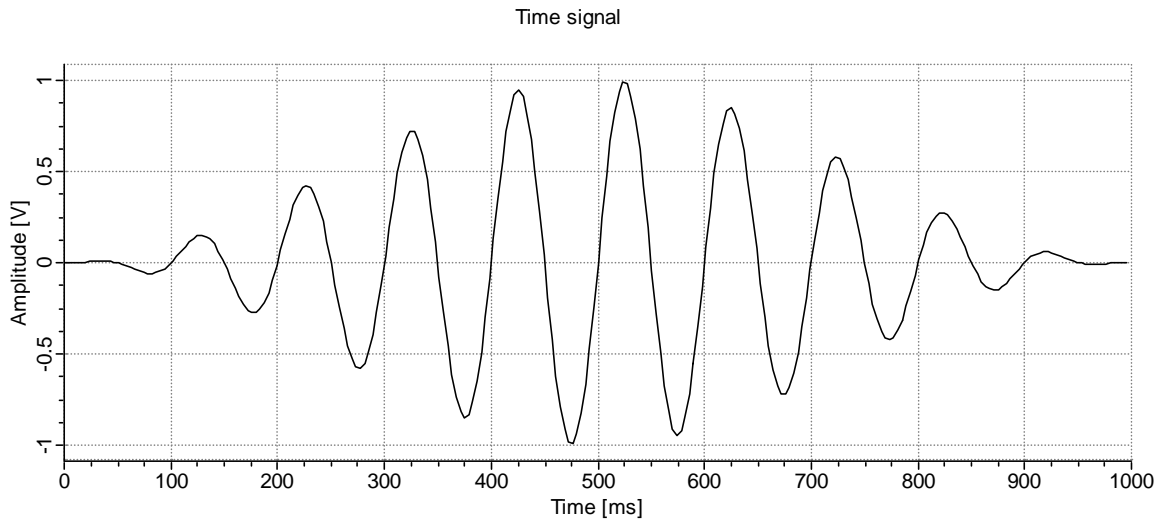
5 Windows

Windows are used to manipulate data (original signal) in such a way, that the desired information can be accurately extracted from the spectrum. In the past, zero padding (frequency spectrum oversampling) was not used much, because it requires more (a lot) memory and fast CPUs. Although memory requirements could have been reduced the CPU speed was still a problem. And some windows were designed to avoid the need of zero padding.

The first thing that one must know about windowing is that often the effect of frequency spill around in the spectral plot, demonstrated in previously, is a direct consequence of the windowing used. The bare fact, that the analyzed signal **is finite in length**, acts like a rectangular Window. We have not applied any transformation, and we already have a window effect. We have already seen how the Rectangular window behaves. Now we can demonstrate some further anomalous matters. Remember. We are looking for a spectrum that will show **the actual** signal frequencies and amplitudes. We hope that windows will get as close as possible and this may depend upon the 'shape' of the window because obviously windows may be differently shaped. Our first example illustrates this by applying a Hanning window to a 10 Hz sine wave.

5.1 Hanning window

The following chart (Picture 14) shows a signal of frequency 10 with applied Hanning window.

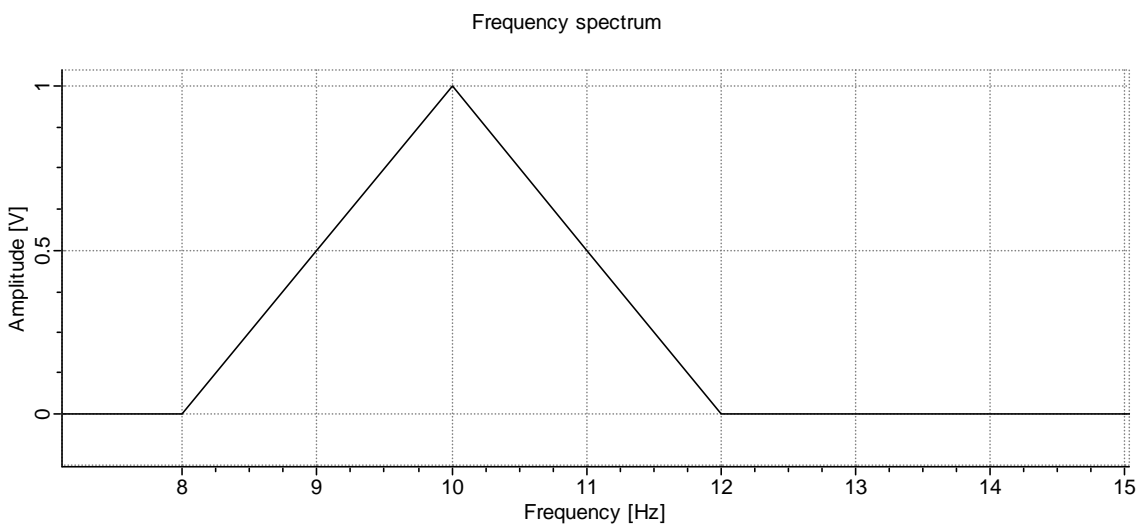


Picture 14 A single tone signal with applied Hanning window

If the array $x[]$ contains the original signal and has n elements then the Hanning window is applied in the following way:

```
for i := 0 to (n-1) do
begin
  c := 0.5*(1 - cos(i*2*Pi/n));
  x[i] := c*x[i];
end;
```

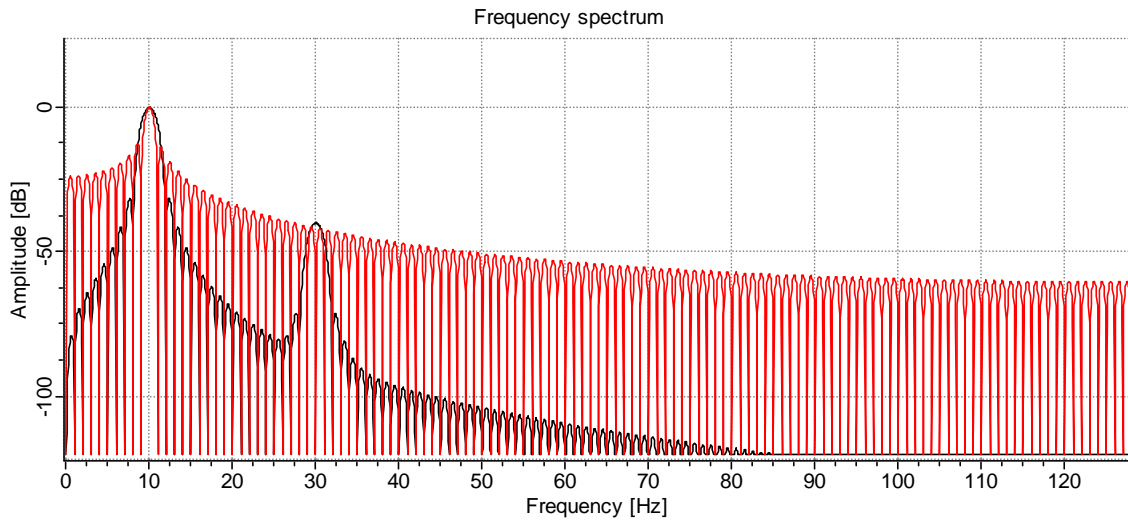
Now let's take a look at how it is with the leakage and amplitude accuracy (Picture 15). The chart shows a zoomed-in peak. If we compare that picture with the zoom from the rectangular window (Picture 3), we see that triangle is 4Hz in width and not just 2Hz as is the case with the Rectangular window.



Picture 15

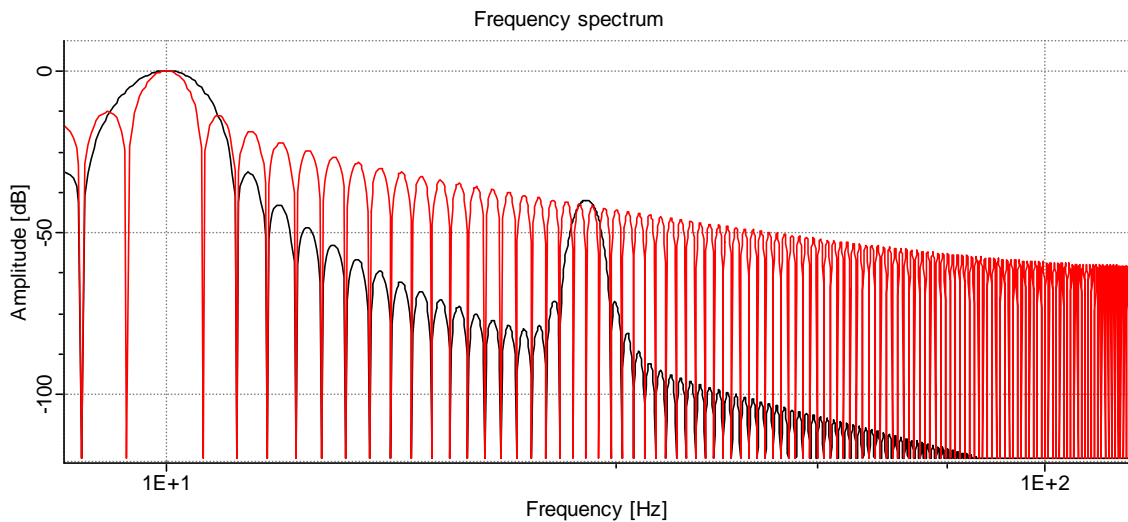
Is this good or bad? It is bad, because now one frequency, even at integer number of periods, pollutes even more spectrum. One problem with leakage is that fact that one strong frequency can obscure the

weaker ones. Let's oversample the spectrum by a factor of 16 (zero padding) and look at it in the logarithmic scale (Picture 16).



Picture 16 Rectangular and Hanning window overlaid

The red series⁶ uses a rectangular window and the black series was set to a Hanning window (FFT Control panel, Window drop down combo). It is clearly seen that the Hanning window affects other frequencies less and that the 30 Hz sine is clearly visible now. Picture 16 shown with logarithmic frequency can be seen on Picture 17.

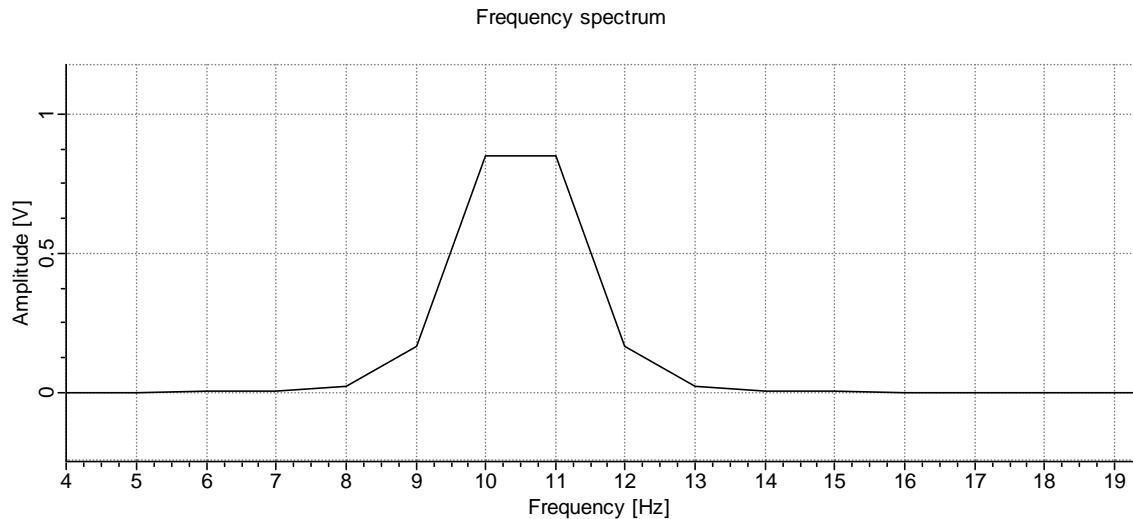


Picture 17 Logarithmic frequency axis

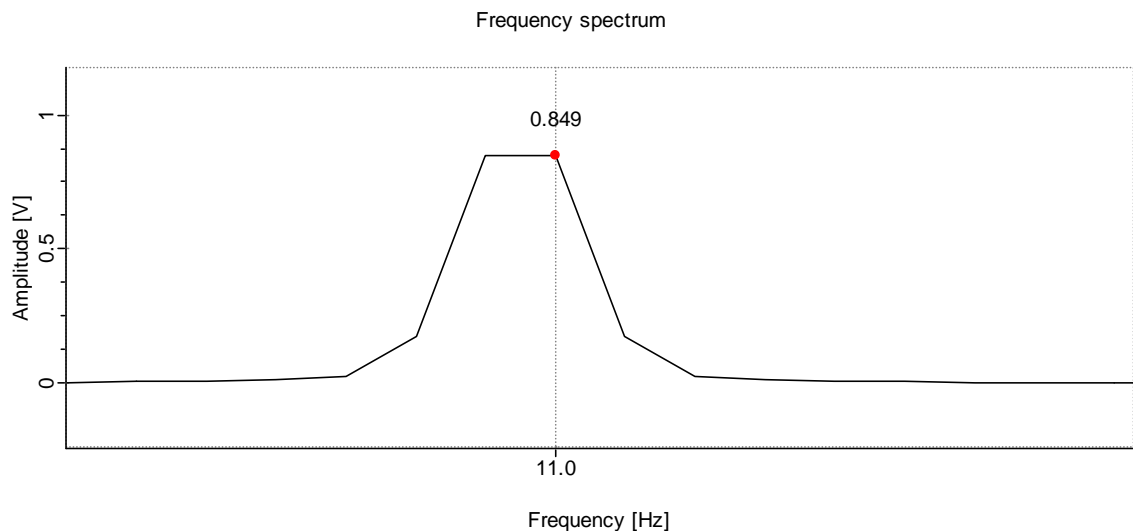
Now we are starting to get an idea of how windows of different kinds may affect the FFT of a signal. The speed of spectral leakage decay (the slope at which spectral leakage falls off) is almost linear with logarithmic frequency and we could talk about dB/decade. And we could talk about the width of the main lobe in percentage of the bandwidth. The rectangular window has a narrower main lobe. Main lobe is the lobe where the frequency is actually present (not leakage).

Let us return to the Hanning window, now, at non-integer number of periods (Picture 18). The frequency of the signal is 10.5 Hz and zero padding is 1 (no oversampling).

⁶ The red line was added to the chart by activating the second channel (check box next to channel 2 on Channels control page). Make sure that Auto scale is active or the display will not be properly aligned. (Right click the spectrum chart to display its popup menu and select -> Peak scale -> Auto)

**Picture 18**

The advantage of the zero padding should now become obvious. If the signal were to be zero padded (spectrum sampled at 16 times more lines) the actual frequency could be considerably more accurately read out, and the same is true of the amplitude, as the next chart indicates (Picture 19). If only the Hanning window is used, the amplitude estimate is a little improved since the biggest error is now only 15% instead of 35% as with Rectangular window and frequency estimate remains the same as without zero padding.

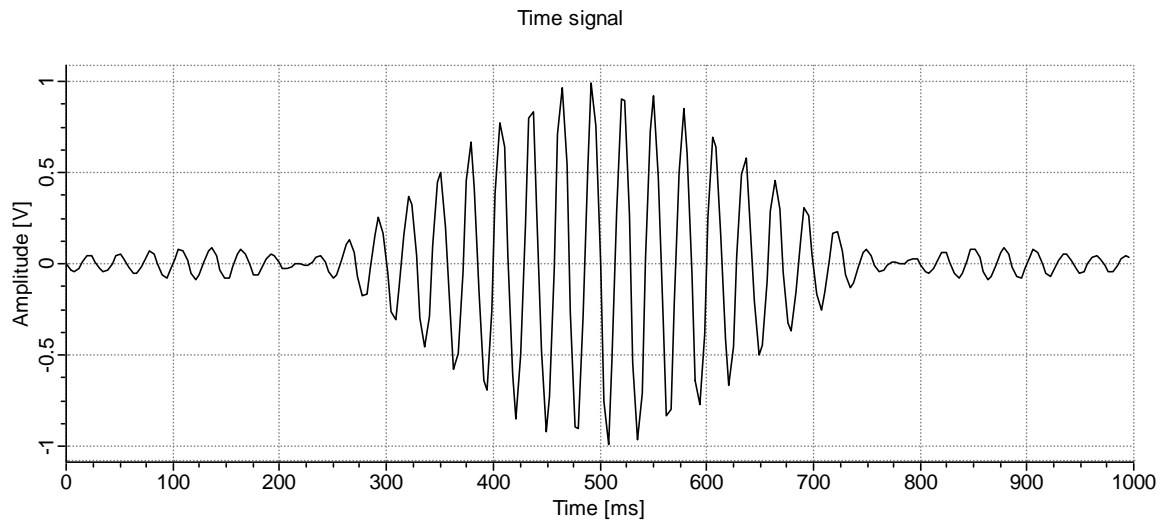
**Picture 19 Effect of Hanning window on a spectral peak**

It is obvious that the leakage effect of the Hanning window is MUCH smaller than with the rectangular window. The only disadvantage of the Hanning window, compared to the rectangular, is the wider peak when we have an integer number of periods. But this is more than compensated for by higher amplitude accuracy and less leakage. We will see later that a wider peak is practically the only way to achieve higher amplitude accuracy, if zero padding is off. There has even been a Window developed especially for that purpose. It is called a Flat top window.

In practice the Hanning window is the most used one. The reason for this is its low leakage and simple formula. In compare to other windows, as we will find out later, Hanning window also offers a good compromise between the width of the main lobe and the attenuation of the leakage.

5.2 Flat top window

Flat top window was especially designed for measurement of the amplitude in the spectrum. But with high spectrum oversampling rates (zero padding) this window has become almost obsolete.

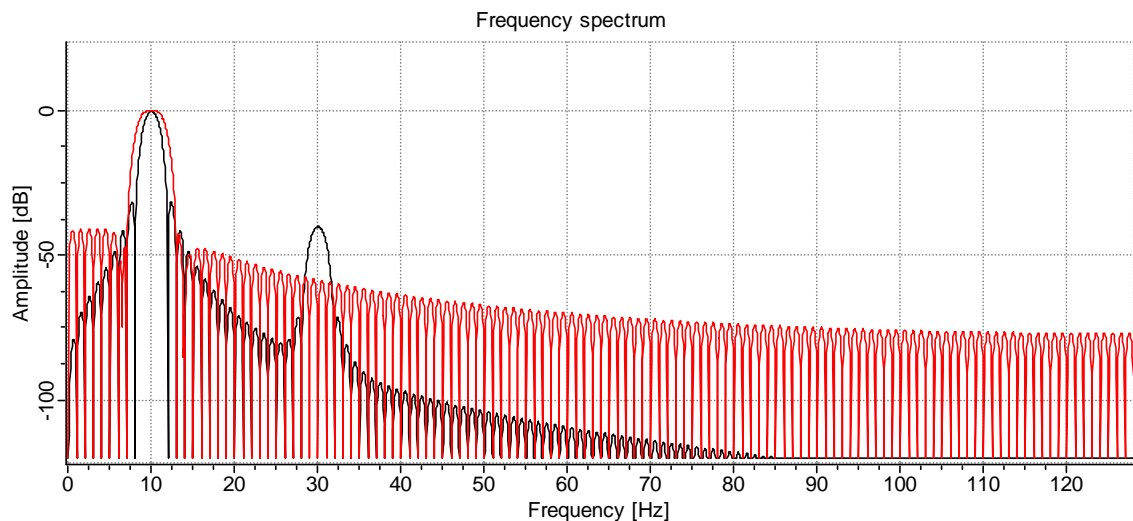


Picture 20 Single tone signal with applied Flat top window.

The window is applied to the original signal as:

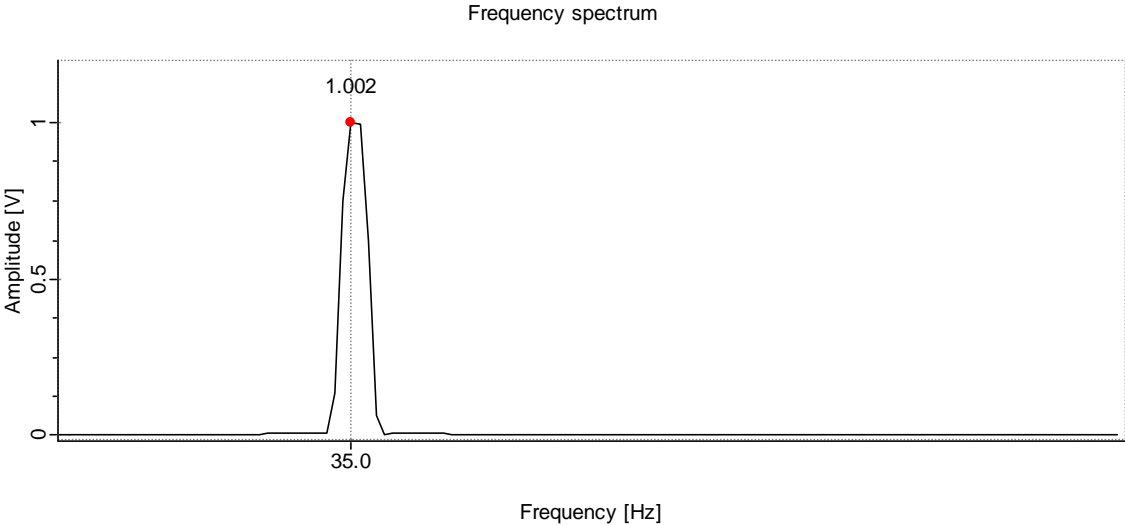
```
for j := 0 to (n-1) do
begin
    w := 2*p*j/n;
    c := 0.2810639 - 0.5208972*cos(w) + 0.1980399*cos(2*w);
    x[j] := c*x[j];
end;
```

The logarithmic spectrum:



Picture 21 Comparison of Flat top and Hanning window on logarithmic scale.

On the Picture 21 the black line is the Hanning and the red line is the Flat Top window. We see that there is less leakage than with the rectangular window and more than with the Hanning window. But then again, this was not the purpose of this window at all. In the worst case the amplitude has only 0.2% numerical error and that without zero padding (Picture 22). The frequency of the signal was 35.4 Hz and its amplitude was 1.



Picture 22 Accuracy of the amplitude estimate with the Flat Top window

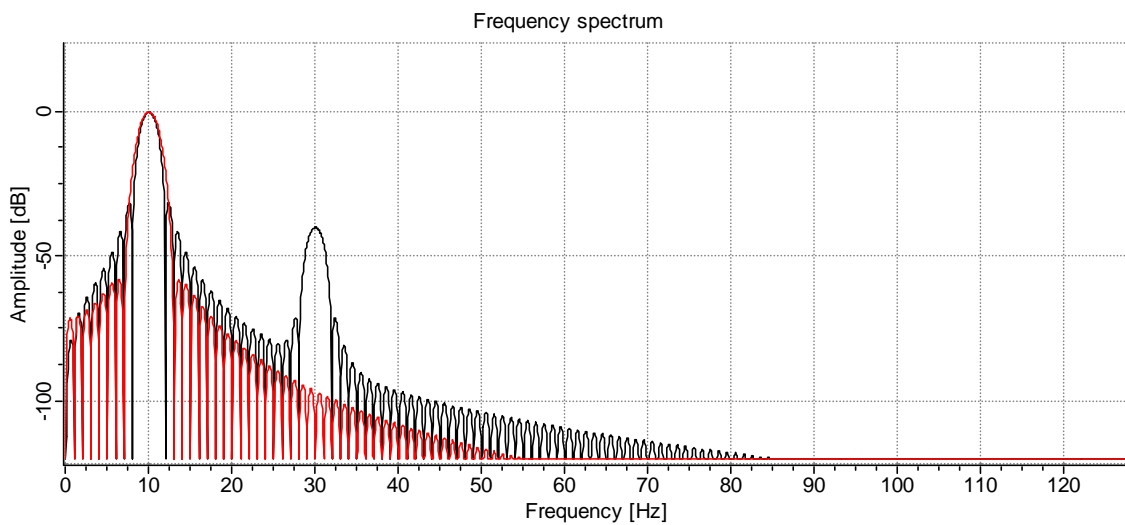
5.3 Blackman window

Since the beginning of frequency analysis in engineering practice, the researchers have tried to find a window that would have the smallest leakage and at the same time a very narrow main lobe. Blackman was one of the first that came really close to that requirement.

It is calculated as:

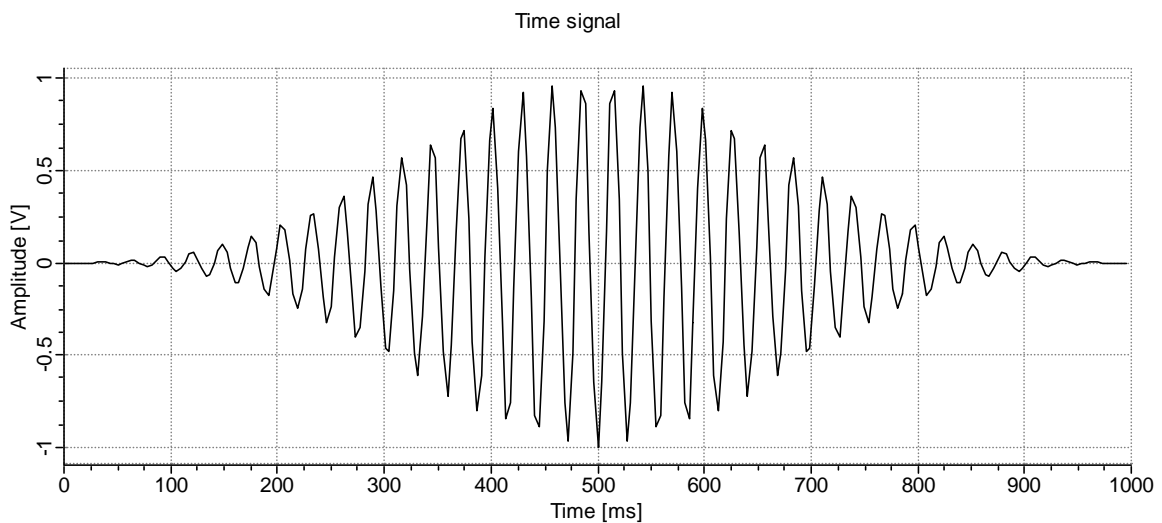
```
for j := 0 to (n-1) do
begin
  c := 0.42 - 0.5*cos(2*Pi*j /n) + 0.08*cos(4*Pi*j /n);
  x[j] := c*x[j];
end;
```

The Blackman window was developed with the search for ideal window. That ideal window is today known as a Kaiser window. Blackman comes very close to the Kaiser window. The Kaiser window was rarely used because it requires the calculation of the Bessel function of the first order.



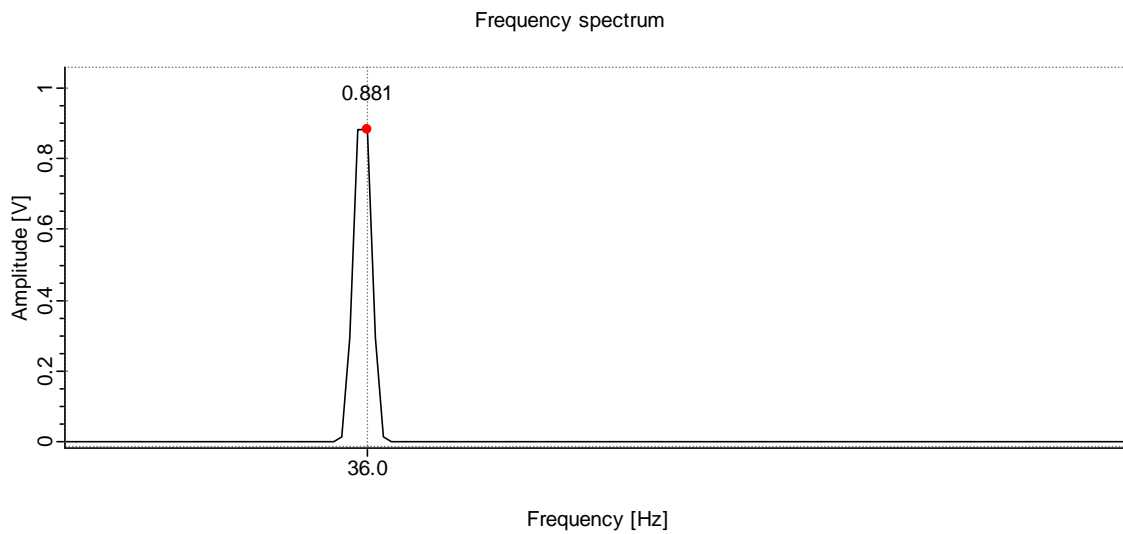
Picture 23 Hanning window (black) and the Blackman window (red)

The Kaiser window is usually not used in today's frequency analyzers. The logarithmic spectrum of a single frequency signal windowed with the Blackman Window is found on Picture 23. The leakage is a little bit smaller than with Hanning window, but the main lobe is even wider.



Picture 24 Blackman window applied to a single tone signal.

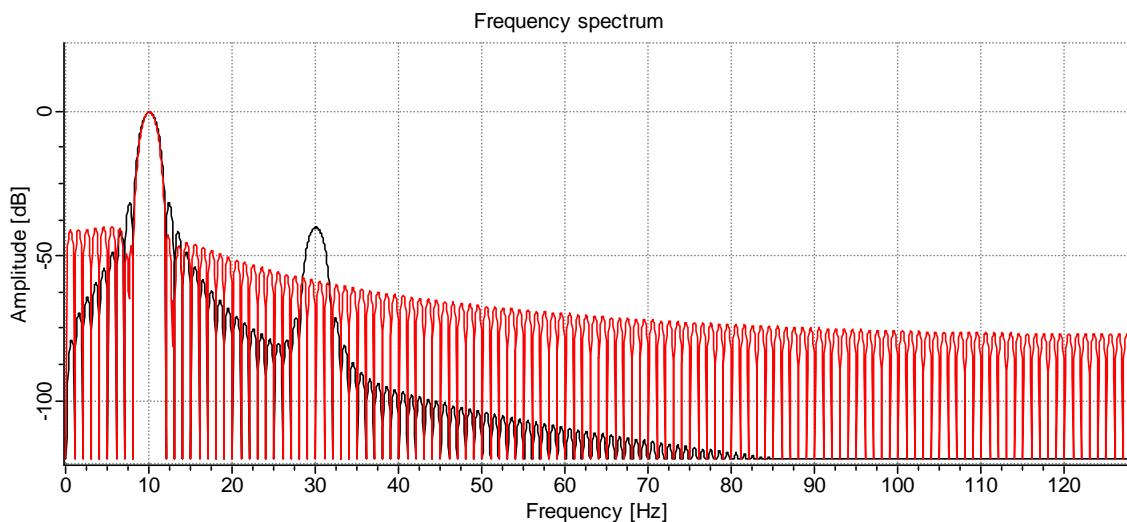
The amplitude accuracy is also a little better than with the Hanning window: 11.9% (with zero padding



off). (The frequency of the signal is 35.5 Hz). The conclusion would therefore be that Blackman window was the best choice for most of the applications, if the Kaiser window was not available and zero padding was not possible.

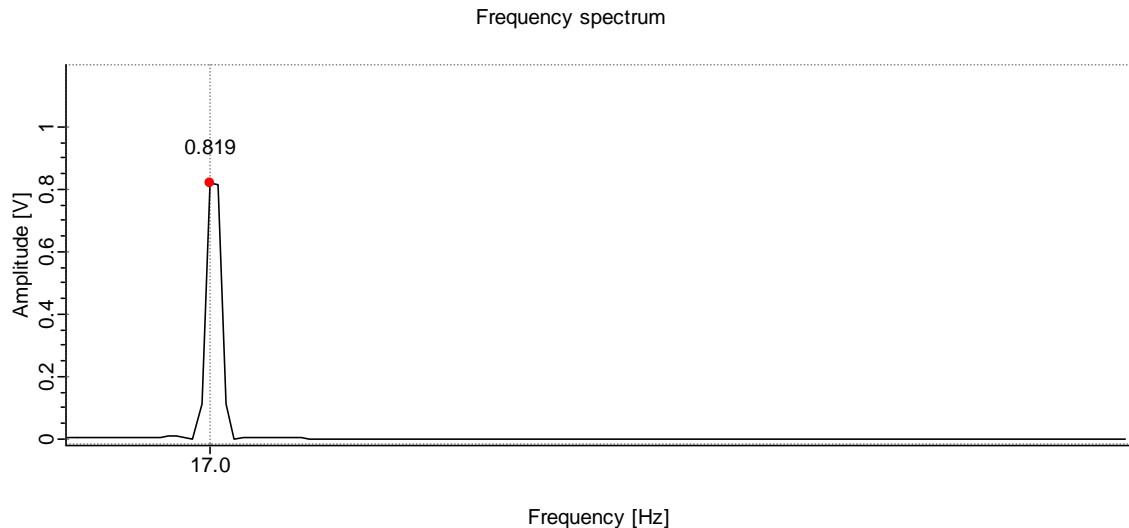
5.4 Hamming window

This window was derived from the Hanning window. The changes were very small but the results are very different. The reason for the changes was complete attenuation of the signal windowed with the Hanning window just at the beginning and just at the end of the signal. After the Hanning window is applied to the signal, those points are almost zero. That means that we lose the frequencies that were present in that area of the signal. The other, much better solution to that problem, was overlapping. When calculating the spectrums, the time signal always contained also the second half of the data contained in the previous time signal part. ⁷ However sometimes we may not arbitrarily long signals available. Let's have a look at the spectrum anyway:



Picture 25 Comparison of Hanning (black) and Hamming window (red)

⁷ The overlapping factor could vary and was usually between 50 and 90%. FFT Properties allows a rich set of overlapping settings on the File input control panel. The Hamming window has lost some of its attractiveness with the procedure.



Picture 26 Maximum error when measuring the amplitude with the Hamming window

We see that the leakage of the hamming window is big (Picture 25), almost as high as with rectangular window. The amplitude accuracy or the maximum error is 18.1% with no zero padding (Picture 26). That is worst than Hanning window and better than with the rectangular.

5.5 Exponential window

When performing dynamic analysis of a system the impulse response of the system must be acquired. As the impulse response decays, the signal to noise ratio is getting worse. For the purpose of lowering the noise in the frequency spectrum, the exponential decay window is applied to the time signal. This window forces the values towards the end of the signal where the signal-to-noise ratio is worst towards zero. What effect has that procedure on the frequency spectrum can be observed by using that window.

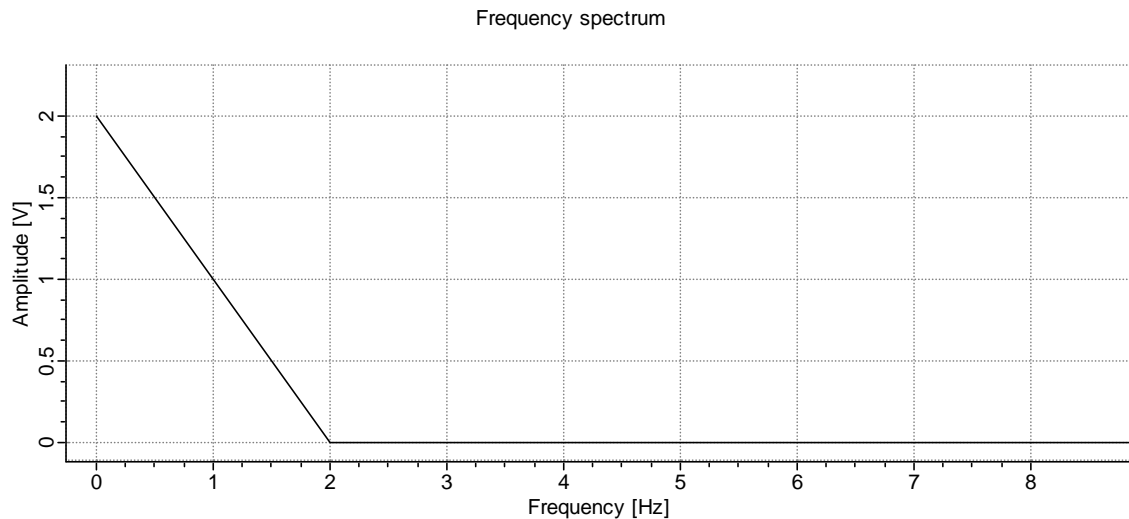
5.6 Summary on windows

To get accurate amplitude and frequency reading for all spectral components you have to attenuate the leakage. By attenuating the leakage you attenuate the effect that all spectral components have on each other. Usually it is sufficient to use the Hanning window, but sometimes it is necessary to use the Kaiser window. For the Kaiser window you can specify how much leakage attenuation you want in dB. But as the leakage attenuation goes up, so does the main lobe grow wider. And by getting wider, you lose frequency resolution. Other windows may prove useful where a special compromise between main lobe width and leakage attenuation is required or the length of the time series is limited.

Once you have selected your window, set zero padding to a factor of at least four and you will get very high accuracy readings with FFT Properties. Too high zero padding factors are also not good, because they substantially increase the processing power required.

6 DC component

If we set DC of signal 1 to 1V, the amplitude of the DC component in the spectrum yields 2 V. This proportion (of 1 : 2) remains the same for all DC offsets (Picture 27). One half of the DC component in the spectrum is exactly equal to the average value of the signal.



Picture 27 The DC component at 0Hz.

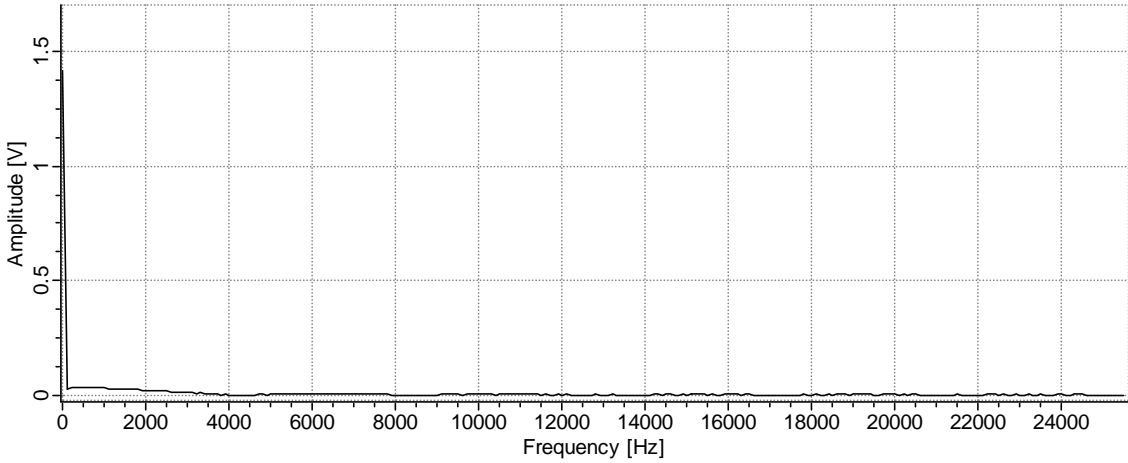
Changing that DC does not have any major effect on the rest of the amplitude spectrum. The DC component in the spectrum resides on the spectrum line 0. The line zero is “reserved” for the DC component and does not tell anything about the frequencies contained in the signal.

When analyzing real world signals the DC component is mostly useless because it doesn't change with time. If the average is changing with time, this indicates, that the length of signal that we are analyzing is too short to resolve the lowest frequencies present in the signal. We know that the highest frequency that we can reliably measure is equal to one half of the sampling frequency. The lowest frequency that we can measure is equal to $1/T$, where T is the length of the signal in seconds. If there are frequencies in the signal lower than $1/T$ [Hz], the DC component will change each time we analyze a different part of the signal.⁸

Many real world signals have at least some very low frequency content. An comparison between display of a signal with and without the DC component can be found on Picture 28 and Picture 29.

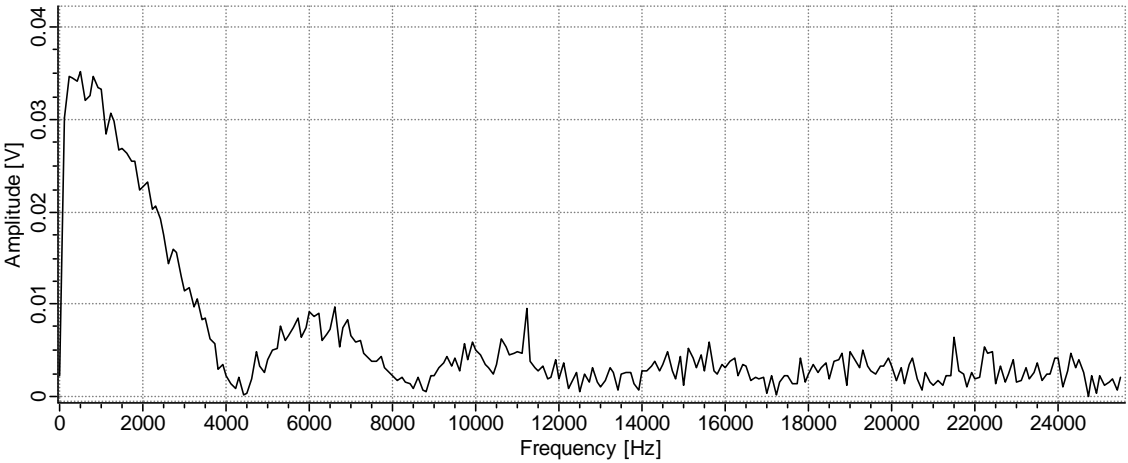
⁸ FFT Properties automatically adjusts the Y axis scale of the spectrum above the maximum and below the minimum of the spectrum. Sometimes the maximum value is the DC component and it may be several times bigger than any other amplitude in the spectrum. In FFT Properties you can check the DC Dump check box and the DC component will be removed without consequence for the rest of the signal. This will help adjust the automatic scales to get a better view of the spectrum. DC Dump is also automatically activated when signals are loaded from files by FFT properties.

Frequency spectrum



Picture 28 DC component present

Frequency spectrum



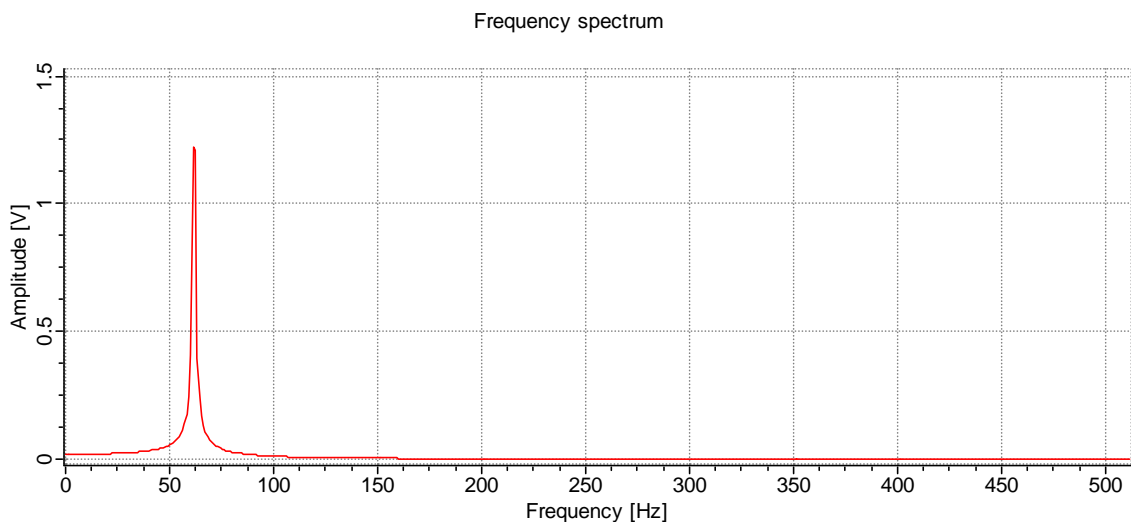
Picture 29 DC component filtered out

7 Beating

FFT has problems dealing with frequencies that are very close together, closer than the width of the peak.

Try the following :

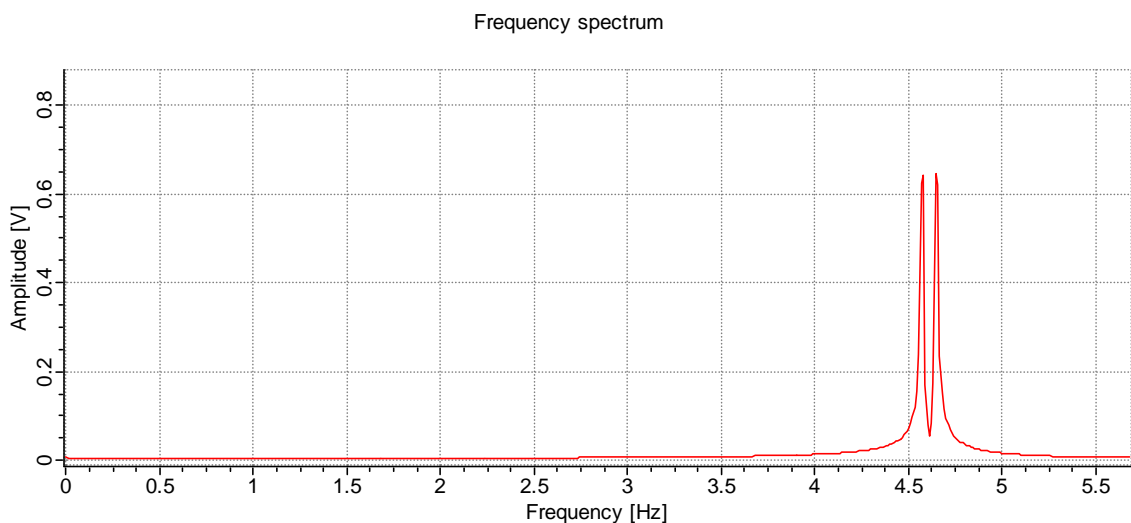
FFT Control panel: Samples = 1024, Sample time = 1000 ms, Rectangular window. Signal generator panel: Signal 1 to $f=61.46$ Hz $A=1$ V, Signal 2 to $f=61.53$ Hz, $A = 1$ V, Harmonic = Off and click the Monitor button. Make sure that Peak hold button is pressed (the S button). The effect is the strongest, if the window is rectangular. What you should see is one single peak that is rising and then falling (Picture 30). This effect is called beating and makes it very difficult to measure the amplitude of both frequencies



Picture 30 Two closely spaced frequencies form a single peak

The problem can partially be solved. The resolution of the frequency spectrum is defined simply by: $1/T$ [Hz], where T is the length of the time signal in seconds (not samples). This is the difference in Hz between two spectral lines. If we can sufficiently increase the length of the signal analyzed, the two merged peaks will resolve.

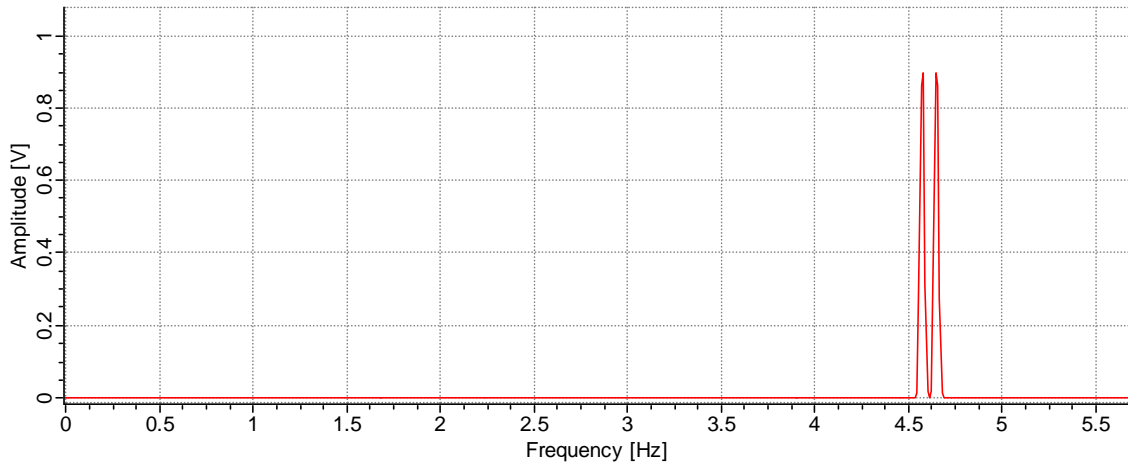
Lets take a look at the same picture (Picture 31) at time signal length of 90s. (90 000 ms).



Picture 31 Two frequencies resolved

The consequence: two frequencies are now resolved. Their amplitude and frequency can be accurately measured if we apply for example the Blackman window (Picture 32) and increase the zero padding⁹. To sum up on the beating: we can get a single peak rising and falling, but there are actually two closely spaced frequencies with constant amplitude. This means that you will get a peak changing its amplitude in dependence of which part of the signal will you be looking at. However, there is another effect that will make the entire spectrum or a part of the spectrum rise and fall, but has nothing to do with beating. That we will examine in the next chapter.

Frequency spectrum



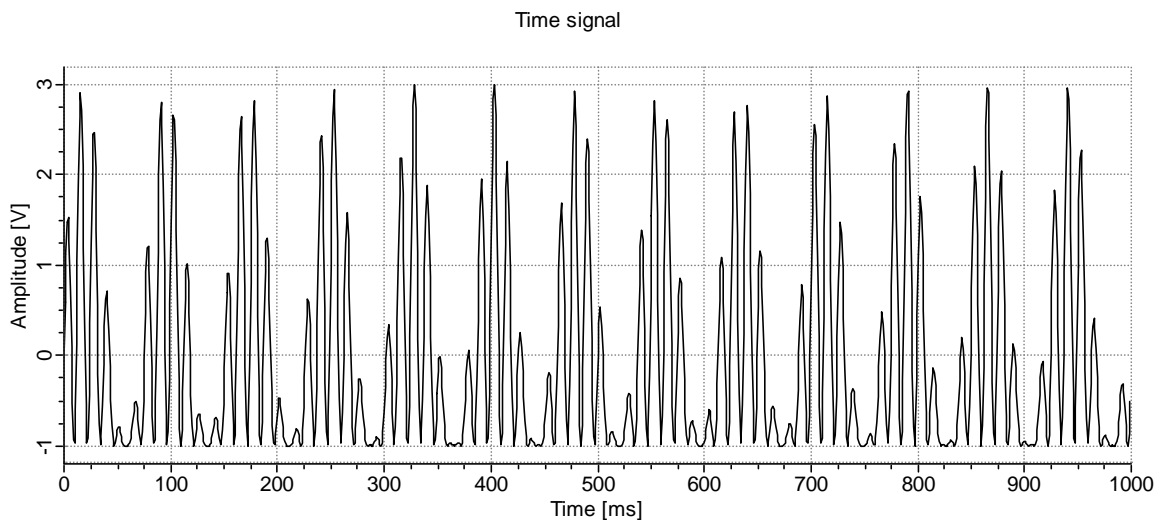
Picture 32

⁹ But not everything is quite as it seems. The frequencies of the two signals seem to be wrong when the two pictures. This is the consequence of aliasing. Both signals were not sampled with the same sampling frequency.

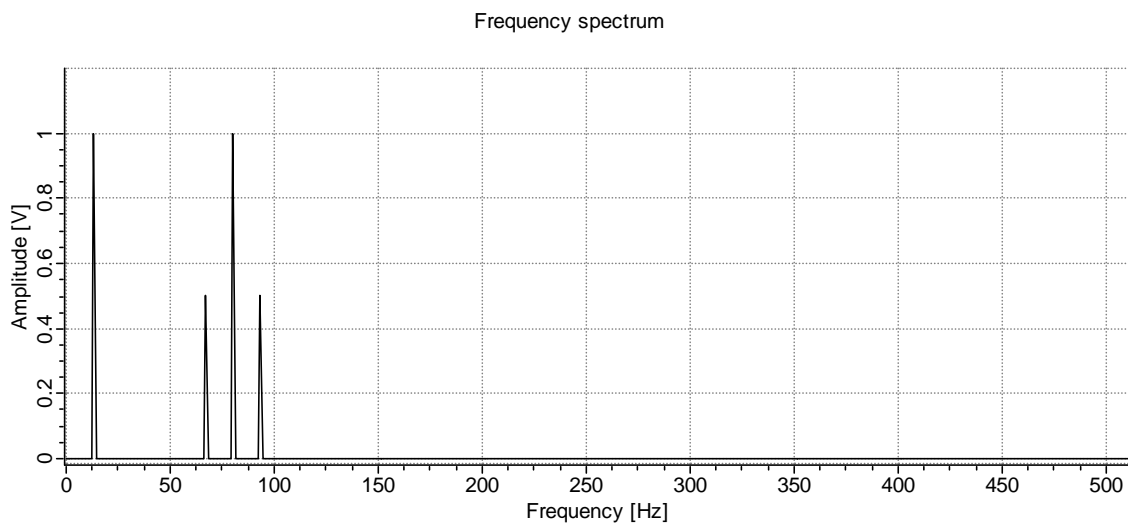
8 Amplitude modulation and sum and difference frequencies

Set signal generator to Multiply Signal1 and Signal2. Make sure that Signal 1 and Signal 2 have values set (amplitude > 0). Now the two signals will be multiplied. Sum and difference frequencies arise in the process called amplitude modulation. (In radio, they are called upper and lower sidebands.) Passing signals through any non-linear (amplitude dependent) process will cause this modulation because such a process has a multiplicative component. Consider this trigonometric identity: $2\cos(a)\cos(b) = \cos(a + b) + \cos(a - b)$. (Similar identities exist for sines alone, and for combinations.) Sum and Difference frequencies can be generated by multiplying the instantaneous amplitudes of two sinusoids.

When interpreting frequency spectrum on rotating machinery such non-linearity's can be quickly detected when a pair of two smaller peaks is found embracing a larger one (Picture 34). This effect is very common and can be found in many real world signals. There are a number of ways in which two (or more) sources can combine. It is not irrelevant, if these mathematical moves can also appear in reality. (Another example: $2\cos^2(a) = \cos(2a) + 1$, This is an example of how a square of a cosine gives a DC component and one frequency at $2x$ higher frequency than the original, but with only half the amplitude.)



Picture 33 Time signal of amplitude modulated signal



Picture 34 Frequency spectrum of amplitude modulated signal

If the frequency of the modulation is very low, we will notice a part or entire frequency spectrum raising and falling.¹⁰

¹⁰ FFT Properties features peak marking tools which allow you to search for Harmonics, Sum and difference frequencies or stand alone peaks.