Introduction to Digital Spectrum Analysis

Peter D. Hiscocks
Syscomp Electronic Design Limited
phiscock@ee.ryerson.ca
www.syscompdesign.com
December 7, 2005

Contents

1 Overview

This paper provides a non-mathematical overview of digital methods of spectrum analysis.

A complex waveform may be treated as being composed of a number of sinusoid waveforms. These sinusoids are of various phases, frequencies and amplitudes. The description of the magnitude, phase and frequency of these various waves is known as the *spectrum* of the signal, by analogy with the spectrum of light.

Spectrum Analysis or *Fourier Analysis* is the process of analysing some time-domain waveform to find its spectrum. We also say that the time domain waveform is converted into a frequency spectrum by means of the *Fourier transform*. This process is reversible: using the *inverse Fourier transform* a spectrum may be converted back into a time-domain waveform.

1.1 Example: Square Wave

Figure 1 shows how a square wave is composed of sine waves of various amplitudes and frequencies.

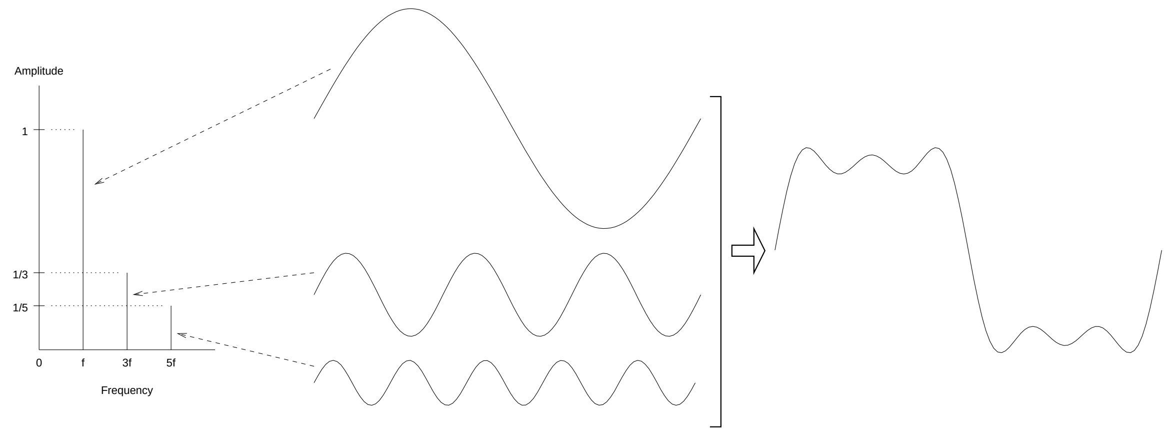

Figure 1: Fourier Series Example

A square wave is composed of sine waves which occur at odd harmonics of the fundamental. As the harmonic number increases, the amplitude of the harmonic decreases.

In the centre of figure 1, we see three of the sine waves that make up the square wave. On the right is the summation of these three sine waves – it is approaching a square wave and would do so more precisely if more harmonics were included.

The left side shows a spectrum display of the waveform. Vertical lines represent the three components of the waveform. The horizontal axis is frequency and the vertical axis amplitude, so the spectrum display shows the frequency at which each harmonic occurs, and the magnitude of the harmonic at that frequency.

The square wave is a special case and in general, one may find harmonics at even multiples as well as odd multiples of the fundamental. As well, there may be cosine waves as well as sine waves present.

2 Applications

Spectrum analysis has a number of applications in electronics and mechanical engineering:

- A pure tone has no harmonics and will show up on a spectrum display as one single vertical line. Distortion of a sine wave will create additional harmonics. Consequently, a measure of the magnitude of the harmonics is a measure of the magnitude of the *harmonic distortion*.

- In a distortion-free (linear) system, two separate input tones (single frequencies) will emerge as the same two tones at the output. If the system is distorting (non-linear), then the system will generate other tones at the sum and difference frequencies of the input signals. A measure of these extra signals is a measure of the *intermodulation distortion*.

- The existence of certain frequencies in a signal may give some clues as to its source. For example, if a signal contains 60Hz, then it is probably picking up interference from the AC power line.

- Power systems frequently manipulate waveforms by chopping them or combining them with other signals. Spectrum analysis allows one to measure the harmonic content of a signal, which may be specified as a requirement.

- The analysis of a mechanical system for resonances can be done by driving the system with a wide-band excitation signal, an impulse hammer blow or random noise from a shaker. Microphones or accelerometers convert the mechanical vibration of the system to an electrical signal. The spectrum analysis of this signal indicates the mechanical resonances in the structure [1].

- The extraction of signals from noise may require some knowledge of the spectrum of the signal and the noise.

- It is useful to see the spectrum diagram for modulation and other signal manipulations.

3 Spectrum Analysis of a Continuous Signal

To understand Fourier analysis, it is useful first to consider the following simple analog spectrum analyser, [2] figure 2.

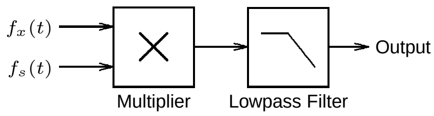

Figure 2: Simple Spectrum Analyser

The unknown signal f_x, which may contain a number of components at various frequencies, is fed into one input of an analog multiplier. A variable-frequency sine wave, the *search frequency* f_s is fed into the other input of the multiplier. The output of the multiplier - which is the product of these two signals - is sent to a lowpass filter and from there to a voltmeter.

To perform a spectrum analysis, the search signal f_s is swept slowly through the frequencies of interest. Multiplying two frequencies together in the multiplier generates two new frequencies at its output, one at the *sum* frequency $f_x + f_s$ and one at the *difference* frequency $f_x - f_s$. The sum frequency is rejected by the lowpass filter. As the two frequencies approach each other, the difference frequency approaches zero. At some point, it falls within the passband of the lowpass filter and appears as an AC signal that can be measured by the voltmeter. If f_x and f_s can be made to coincide, then the output from the lowpass filter is a DC signal proportional to their magnitudes.

Consequently, spectrum analysis can be performed by sweeping the oscillator frequency f_s and noting both the frequency and amplitude of the signal whenever a signal appears at the output of the lowpass filter.

There are several interesting aspects of this system:

- The system may be viewed as one using a *swept bandpass filter*. The centre frequency of the filter is established by the search frequency. The bandwidth of the analysis filter is twice the cutoff frequency of the lowpass filter.

- Alternatively, the system may be viewed as a *correlator*. Correlation is a measure of the similarity of two signals, and may be obtained by multiplying two signals and averaging the result over a period of time. In this system, the correlation takes place between sine waves in the unknown signal and the sine wave search frequency.

- If the search frequency could be made *exactly* coincident in frequency to some sinusoidal component of the unknown, then the output voltage would be a DC value proportional to the phase angle between these two sine wave signals. In practice, it is difficult to synchronize the two signals, and the output of the lowpass filter is the difference frequency between the unknown and search frequency. This difference frequency only approaches zero (DC).

3.1 Fourier Series

Mathematically, a periodic signal x(t) may be represented by a Fourier Series as follows:

$$x(t) = A_o + \sum_{k=1}^{\infty} B_k \cos k\omega_o t + \sum_{k=1}^{\infty} C_k \sin k\omega_o t \tag{1}$$

where the variables are:

A_o	the *average* (DC) value of the signal: $A_o = \frac{1}{T}\int_0^T x(t)dt$
k	the *harmonic number*: 1=fundamental, 2=2nd harmonic, etc
B_k	peak value of the magnitude of the kth cosine harmonic
C_k	peak value of the magnitude of the kth sine harmonic
ω_o	*fundamental* frequency, $\omega_o = \frac{2\pi}{T}$
T	period of $x(t)$

The way to read this is as follows:

- The original waveform $x(t)$ is some repeating function of time, such as a square wave or triangle wave.
- This waveform repeats every T seconds, where T is the *period*.
- The waveform may be constructed by adding together a DC term plus cosine and sine waves of various amplitudes.
- The DC term is the average value, which is found by integrating over one complete cycle of $x(t)$and dividing by the period.
- The cosine and sine waves occur at multiples of the fundamental frequency $\omega_o = 2\pi f_o$, $f_o = 1/T$. In other words, the fundamental is the lowest complete sine wave that can be fitted into $x(t)$. The sine waveforms are referred to as the *inphase* components. The cosine waveforms are referred to as the *quadrature* components.

  Depending on the shape of $x(t)$, the quadrature waveforms may be present, or the inphase waveforms, or both.
- Not all multiples of the fundamental need be present. For example, a square wave includes only the odd-numbered multiples.
- The terms B_k and C_k determine the magnitude of each harmonic. These are the quantities that show up as vertical lines on a spectrum plot.

3.2 Magnitudes of the Harmonics

It may be shown that the magnitudes B_k and C_k of the harmonics are given by:

$$B_k = \frac{2}{T}\int_0^T x(t)\cos(k\omega_o t)dt \tag{2}$$

and

$$C_k = \frac{2}{T}\int_0^T x(t)\sin(k\omega_o t)dt \tag{3}$$

For example, to find the magnitude B_k of the quadrature component:

- multiply the original waveform xt by a cosinusoidal waveform of frequency k times the fundamental frequency ω_o
- integrate the result over a complete cycle of the waveform
- double the result

To the extent that lowpass filtering approximates integration (which is approximately true when the frequency is above the corner frequency of the lowpass filter), this description is exactly as we described for figure 2. Alternatively, this *multiply and integrate over one cycle* operation may be regarded as correlation between the unknown signal $x(t)$ and a sine wave of frequency $k\omega_o$.

Similarily, to find the magnitude C_k of the inphase component, we multiply by a *sinusoidal* waveform and integrate over a complete cycle.

In our orginal spectrum analyser design, we correlated against one sine wave, without worrying about phase. That was acceptable because the search frequency will never be exactly equal to the frequency in $x(t)$, and so the phase relationship between them will be continuously increasing or decreasing.

Reiterating: equation 1 indicates that at any given harmonic frequency of the fundamental, there may be a quadrature (cosine) and inphase (sine) waveform, that is, the harmonic contains both a sine and cosine wave.

Consequently, a spectrum analysis should have two components: one spectrum showing the magnitude of the various sine (inphase) components, and one spectrum showing the cosine (quadrature) components. These components may be regarded as two vectors at 90^o to each other. The magnitude of these vectors can be shown in *rectangular* format or in *polar* (magnitude and phase) format.

The inphase/quadrature presentation is difficult to interpret and the phase of a spectrum is often ignored. Consequently, of these possible ways of displaying the harmonic vectors, the most useful is the polar-magnitude presentation.

Figure 3 shows how the inphase and quadrature components are determined. The search frequency is generated by a *quadrature oscillator*, which produces a sine wave and cosine wave at various frequencies.

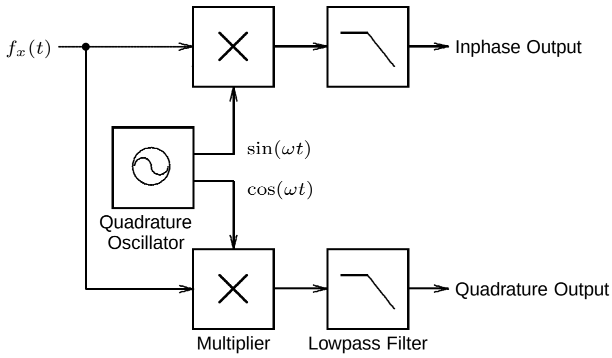

Figure 3: Simple Spectrum Analyser

Figure 4 shows additional stages in which the inphase and quadrature components are converted into polar format.

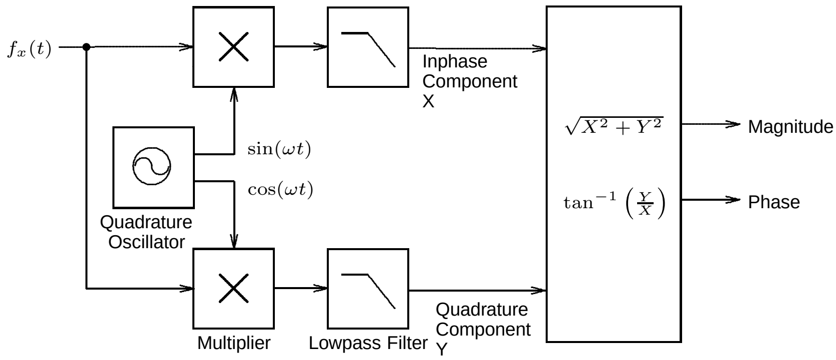

Figure 4: Magnitude and Phase Format

4 Frequency Analysis using Digital Techniques

A periodic signal $x(t)$ may be represented in digital form as a series of numerical values. One cycle of this waveform data can be stored in a computer and then operated on according to equations 1, 2 and 3 to determine the harmonic content of the waveform.

This process is very convenient and makes possible certain signal processing operations that would be costly and difficult in analog circuitry. However, converting the signal from the *analog domain* to the *digital domain* introduces some limitations and inaccuracies.

First we'll show how the spectrum analysis is accomplised, and then in section 4.5 and following we'll discuss some of the limitations. The ideas in this section are based on the very clear exposition in reference [7].

4.1 Discrete Sine Transform

Consider that the instrument captures a waveform record that is 16 samples long. By analogy with the analog spectrum analyser described previously, we can determine the presence of a frequency in the record by multiplying it by the sinusoid of that some *search* frequency and integrating the result.

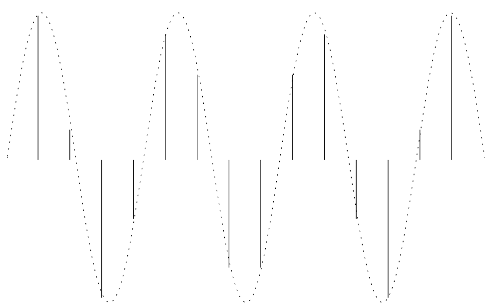

Figure 5: Sampled Sine Wave

In this case, we multiply each of the 16 points on the waveform record by the corresponding 16 points on the sampled version of the search frequency sinusoid. For example, a search frequency of 3.5 times the fundamental is shown with its 16 samples in figure 5. (The first and last samples are zero.)

We sum up each of these products and then we divide this sum by the number of samples to normalize it. The result is the amount of that frequency present in the original waveform.

We repeat this operation for each of the possible frequencies in the spectrum. Finally, we present the results as a spectrum display.

4.2 Frequency Range and Resolution

Based on this concept, we can now explore the issues of resolution and frequency range.

Some definitions:

- N is the total number of samples in the waveform record.
- k is the index number of a sample, between 0 and $N - 1$.
- Δt is the time increment between each sample in the record.
- T is period of some sine wave search frequency, and $f = 1/T$ is the frequency of that same sine wave frequency.

The largest period T_{max}, or *lowest* search frequency f_{min} is one that spans a half cycle in the same interval of the waveform record. Figure 6 shows the lowest possible search frequency spanning the record length, with N=16 samples, each of interval Δt.

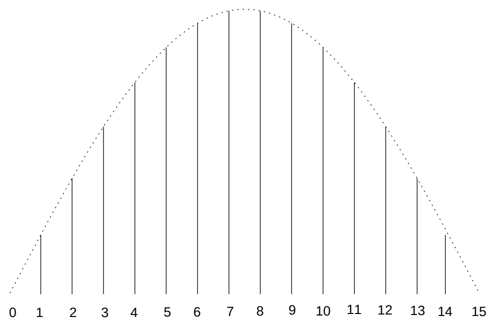

Figure 6: Lowest Search Frequency

We can relate the period of this search frequency to the number of samples and time increment as follows:

$$\frac{T_{max}}{2} = N\Delta t \tag{4}$$

or

$$f_{min} = \frac{1}{2N\Delta t} \tag{5}$$

Notice that the lowest analysis frequency decreases with increasing record length.

Example

For an instrument that captures 500 samples at 10msec intervals, what is the lowest frequency that can be resolved?

Answer

$$
\begin{aligned}
f_{min} &= \frac{1}{2N\Delta t} \\
&= \frac{1}{2 \times 500 \times (10 \times 10^{-3})} \\
&= 0.1 \text{ Hz}
\end{aligned}
$$

The smallest period T_{min} or *highest* search frequency f_{max} is one for which one-half cycle spans two samples, as shown in figure 7 for 16 samples[1].

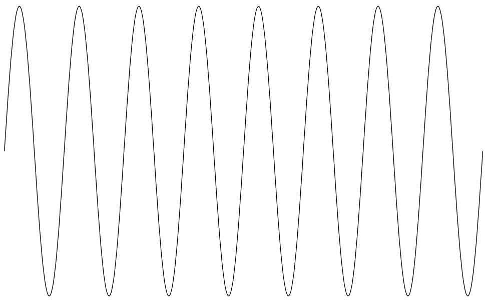

Figure 7: Highest Search Frequency

Δt is the interval between each sample (waveform zero crossing, in this case), so that one complete cycle of this search frequency corresponds to two sample intervals. Then:

$$T_{min} = 2\Delta t \tag{6}$$

or

$$f_{max} = \frac{1}{2\Delta t} \tag{7}$$

Example

For the previous instrument, what is the maximum frequency shown in the spectrum?

Answer

$$
\begin{aligned}
f_{max} &= \frac{1}{2\Delta t} \\
&= \frac{1}{2 \times (10 \times 10^{-3})} \\
&= 50 \text{ Hz}
\end{aligned}
$$

The number of spectrum lines is the ratio of maximum to minimum frequency, wwhich is simply N, the number of samples.

In the example, the spectrum would contain 500 frequencies between 0.1Hz and 50Hz.

[1]Each sample point coincides with a zero crossing, so all the samples of this particular waveform are zero. That's normal and correct.

4.3 Discrete Fourier Transform (DFT)

As we explained previously, the input waveform will in general contain frequencies of various phase angles – they need not be sine waves. To detect a frequency component of some arbitrary phase, as indicated in the analog spectrum analyser of figure 3, it is necessary to use both *inphase* and *quadrature* search frequencies.

Then the magnitude of frequency component can be obtained applying the hypotenuse formula to the inphase and quadrature components:

$$Magnitude = \sqrt{inphase^2 + quadrature^2} \tag{8}$$

This is the most used result of a spectrum analysis. In some situations, the phase of the spectral component is also useful:

$$Phase = \tan^{-1}\left(\frac{quadrature}{inphase}\right) \tag{9}$$

There is one further wrinkle to the DFT, and that is the matter of *negative frequencies*.

We normally think of a sine wave as the result of a vector that rotates in a clockwise direction. However, a given sine wave can also be visualized as the sum of two vectors, one rotating clockwise, the other counterclockwise (figure 8). The vectors are equal in length, each is half of the total. The clockwise rotating vector may be regarded as a positive frequency. The counterclockwise rotating vector is a negative frequency.

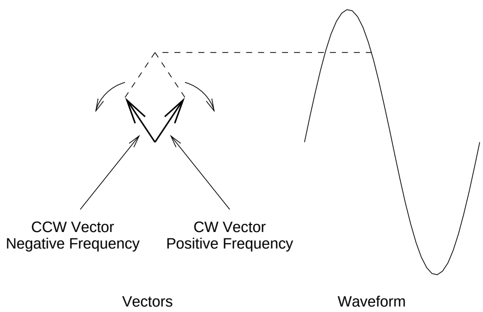

Figure 8: Positive and Negative Frequencies

Normally, this is irrelevant, but it shows up in the DFT spectrum by creating two sets of frequencies: one positive and one negative. These spectra are mirror images of each other, so one of them (usually the negative spectrum) can be ignored if the amplitude of the other one is doubled[2].

As in the case of the discrete sine transform, there are N frequencies in the complete spectrum, where N is the number of samples in the waveform record. Half these points are in the negative frequency spectrum and half in the positive spectrum. Consequently, there are $N/2$ useful frequency points in the positive spectrum.

4.4 Fast Fourier Transform

Assuming that the search frequency sine waves are pre-computed, the discrete fourier transform requires in the order of N^2 multiplications. Many of these multiplications can be simplified. For example, multiplication by zero produces a zero result. Multiplication by -1 changes the sign of the operand.

[2] A more detailed and mathematically credible explanation is given in reference [5].

As well, because the sine and cosine functions are periodic and repeat themselves, many of the multiplications are redundant. Once done, the product can be used repeatedly. The *fast fourier transform* (FFT) takes all these shortcuts into consideration and reduces the number of multiplications from N^2 to $1/2N\log_2(N)$. So a 512 point transform would reduce from $512^2 = 262144$ multiplications to $0.5 \times 512 \times 9 = 2304$ multiplications, a factor of 114. (To some extent this speedup is offset by the necessity of shuffling the input or output data into a particular order, but regardless the FFT will be much faster than the DFT.)

The FFT requires that the record length be a power of 2: 512, 1024, 2048 ... data points.

For transient signals, the data record dies away to zero and consequently can be padded with zeros up to a length that is the next highest power of 2. For example, a 500 point record could be padded with 12 zeros to increase it to 512 points.

A periodic signal could be windowed in such a way that it tapers down to zero at the beginning and end of the record, and then padded with zeros to bring it up to the required length. However, this may affect the frequency content of the spectrum, so it must be done with care.

In the next section, we discuss some of the limitations of digital calculation of the DFT or FFT.

4.5 Dynamic Range

The *dynamic range* of the spectrum analysis determines the smallest spectral component that can be observed. That is, it determines the *noise floor* of the measurement. A signal that is below the noise floor cannot be observed. This could be important when looking for small levels of distortion, which would show up as low-amplitude harmonics of the input signal.

Using analog techniques, such as those described previously, the noise floor is determined by the thermal noise of the resistances and the noise figures of the various semiconductor devices. These noise signals are small, and the effect of narrowband filtering (which is the essence of spectrum analysis) is to reduce the noise even further. Consequently, it is relatively straightforward to obtain a dynamic range in the order of 80db.

Using digital techniques, an analog signal is converted to digital numbers with a precision that is determined by the number of bits in the analog-digital (A/D) converter. For example, a ramp signal is represented by a series of discrete values.

Figure 9 illustrates the process with a 3-bit analog-digital conversion followed immediately by 3-bit digital-analog conversion.

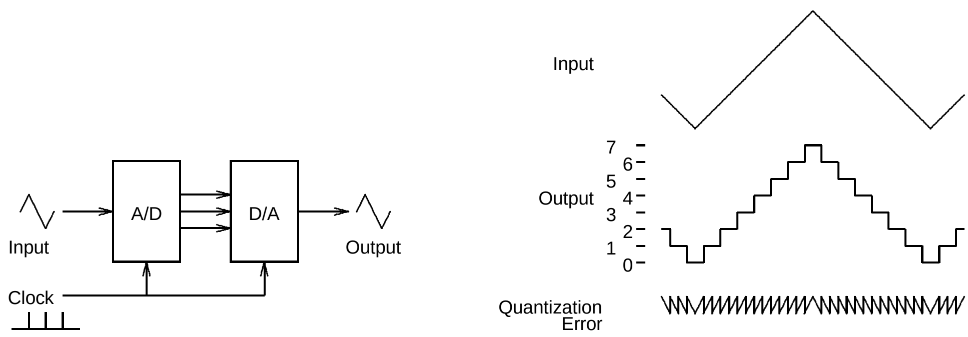

(a) Circuit (b) Waveforms

Figure 9: Signal-Noise Ratio

In effect, the smooth ramp is converted into a staircase. Each additional bit doubles the resolution of the conversion, or, what is the same thing, halves the staircase step size.

The *quantization error* is the difference between the input and output signal. The quantization error may be regarded as a type of noise. It may be shown (see Johns and Martin, [4]) that the RMS value of the quantization noise is given by

$$V_{noise} = \frac{V_{lsb}}{\sqrt{12}} \qquad (10)$$

where V_{noise} is the RMS value of the noise and V_{lsb} is the input voltage equivalent to one step on the converter. This noise value is fixed by the number of bits and the reference voltage of the converter. Consequently, for any given input signal magnitude, it would be possible to use equation 10 to determine the signal-noise ratio.

Ideally, the signal should be much larger than this noise level. In logarithmic terms, each additional bit approximately increases the signal-to-quantization-noise ratio by 6db. For example, an 8 bit converter (which is common in digital oscilloscopes) will provide a 48db signal-noise ratio. This is the best possible case, since it assumes the input signal drives the A/D converter to its maximum possible output, but does not overload it.

4.6 Aliasing

A Mechanical Example

The effect of aliasing can be illustrated by an example from mechanical engineering. Consider the situation where a rotating disk is illuminated with a stroboscope[3].

If the strobe flashes several times during each revolution of the disk, we will see it as a series of consecutive visual samples, and the samples will show the disk rotating in the correct direction. In this case, we say that the sampling frequency is high compared to the frequency of the disk.

Now consider that the strobe frequency is reduced so that it flashes once per revolution of the disk. This time, all the samples are taken with the disk in the same position, so it appears to be stopped.

If the strobe frequency is reduced slightly from the previous value, the disk will appear to be rotating slowly in the *backwards* direction. This is a visual illusion, and the effect is known as *aliasing*.

Notice that the apparent speed of rotation of the disk is equal to the difference between the actual rotational speed of the disk and the stroboscope sampling frequency. Put another way, this system creates a visual frequency at the difference between the rotational frequency and the sampling frequency.

Sampled Electrical Signal

When a signal is captured by a digital oscilloscope, it is converted into a series of samples, each one a digital number. Analogous to the stroboscope example, this sampling process creates a difference frequency. Under some circumstances, this difference frequency can be useful[4]. However, in spectrum analysis, the difference frequency is a potential nuisance.

Consider for example that some frequency f_x is sampled at a rate f_s, where $f_s = 1.3 f_x$. Then a signal will appear at a frequency of $f_b = 1.3 f_x - f_x = 0.3 f_x$. This frequency was not in the original, it is an artifact of the sampling process.

There are two steps to the solution to this problem.

- Ensure that the sample frequency is greater than twice the highest value of the unknown frequency. Any difference frequencies will then lie *above* the highest value of the unknown frequency.
- Suppress the display of frequencies above the unknown frequency, since they will contain aliases.

[3] A *stroboscope* or *strobe lamp* is a flash lamp driven from a variable frequency oscillator. It produces a bright flash of light at the specified interval.

[4] This is the princible behind *equivalent time sampling* which can be used to make a repetitive signal appear as a lower-frequency replica of the original.

Put another way, one must ensure that there are no frequency components in the unknown signal that are more than one-half the sampling frequency. This requirement is often enforced by an *anti-aliasing* lowpass filter which restricts the maximum frequency of the input signal. This is not difficult to implement when the sampling frequency is fixed. However, in a digital oscilloscope, the sampling frequency necessarily varies over a wide range which requires that the anti-aliasing filter have a variable cutoff frequency. This turns out to be difficult to implement in a low-cost oscilloscope. Instead, the operator must learn to operate the scope in such a way that alias frequencies can be prevented or identified as such and ignored.

4.7 Leakage

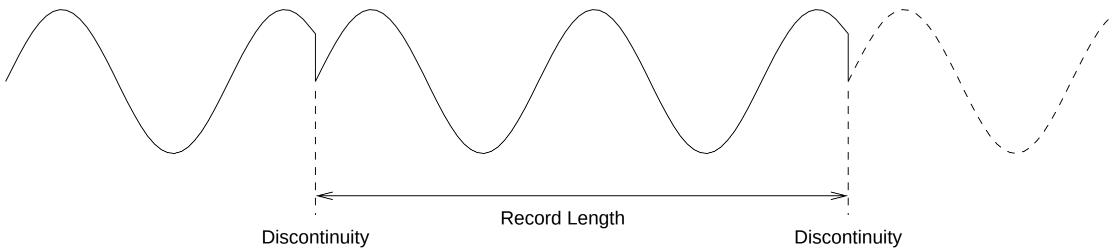

Figure 10: Truncated Record

The oscilloscope works in a batch mode. It collects a group of samples, which are collectively known as a *record*, and then displays them directly as a waveform display. This *capture-then-display* process happens at a rate that depends on the timebase setting, but is typically 5 records per second.

When spectrum analysis is occurring, these waveform records are handed off to the analysis software, which then extracts the various component frequencies and displays them as a spectrum.

In general, one cannot expect that a complete integer number of cycles of the input waveform will be occur in each waveform record. Consequently, there are discontinuities in the waveform where one record ends and the next one starts (figure 10). From the point of view of the spectrum analyser, the discontinuities at the beginning and end of each waveform are part of the signal. These discontinuities show up in the spectral diagram as *leakage*, that is, unwanted components in the spectral diagram. The correct components of the spectrum appear to leak into other frequencies.

The effect of these discontinuities is more evident for shorter records, where the discontinuity is a larger part of the record.

The discontinuities occur at the beginning and end of each record, and so their effect may be minimized by a *tapering function* or *window function*. A window function has two effects: it reduces the spurious frequencies created by record discontinuities, and it broadens the display of individual frequency components. There are tradeoffs between these two effects and these have given rise to various window functions.

Triangular Window

The process of windowing a waveform is shown in figure 11.

Figure 11(a) shows a waveform record that is truncated part way through the fifth cycle. Figure 11(b) shows the triangle windowing function. It starts at zero, runs up to unity and then back down to zero, thereby depreciating the start and finish of the record.

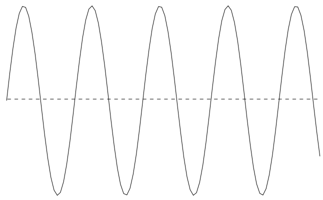

(a) Original Waveform

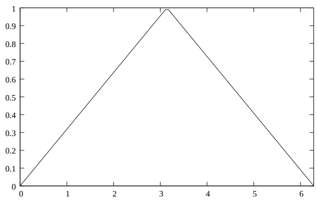

(b) Triangle Window

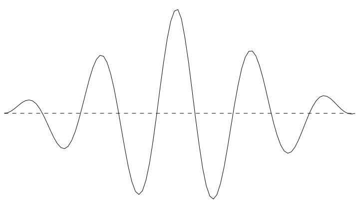

(c) Windowed Waveform

Figure 11: Triangle Window

The value of the triangular window function is:

$$w(n) = 1 - \frac{|2n - N + 1|}{N} \tag{11}$$

where $w(n)$ is the value of the triangular window at sample n, the sample number runs from 0 to $N - 1$, and there are N samples in the record.

Figure 11(c) shows the windowed waveform, which is the product of the original waveform in figure 11(a) times the windowing function in figure 11(b). Notice that there are no discontinuities at either end of the windowed waveform.

Hamming Window

The *Hamming* window, another common window, is shown in figure 12.

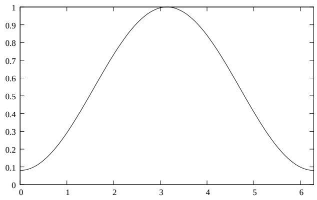

Figure 12: Hamming Window

The value of the Hamming window function is:

$$w(n) = 0.54 + 0.46 \cos \left(\frac{[2n - N + 1]\pi}{N} \right) \tag{12}$$

A number of other window functions are given in reference [6].

Self-Windowed Waveform

A window function is *not* appropriate for a waveform that is the result of a transient. For example, the result of applying an impulse to a system could be a sine wave that starts at a large amplitude and then dies away (ringing)[5]. The discontinuity at the beginning is not an artifact of the system, it is real. As well, the waveform dies away on its own accord, it does not need to be tapered. Consequently, applying a window function to this waveform is not necessary since it is in some sense self-windowed.

4.8 Picket Fence Effect

The spectrum analyser created by the DFT is in effect a group of bandpass filters centred at the analysis width Δf. The shape of these bandpass filters[6] combined is such that the overall spectrum response is uneven, consisting of a series of bumps. If a pure tone is swept through these filters, it will vary in amplitude as it coincides with the peak at the centre of a bandpass filter and the valley between two filters. This is the *picket fence effect*.

To obtain an accurate measurement of the amplitude of some spectrum component, it is necessary to ensure either that it is centred in a bandpass filter or that the overall response of the filters, combined, is acceptably uniform.

4.9 Other Methods of Spectrum Analysis

The DFT and FFT synthesize spectrum analyser that contains N bandpass filters (often referred to as *bins*). The width of these bandpass filters in Hertz is fixed. We refer to these filters as *constant bandwidth*. This type of analysis is most appropriate for analysing a narrow-band frequency such as as a resonance in a mechanical structure.

In audio work, it transpires that this type of analysis is best performed with filters that are *constant percentage bandwidth*. These filters are typically set to pass such one octave, one-third octave or 1/12 octave. The energy

[5]This is a common technique for determining the resonant frequencies of a mechanical structure. The structure is hit with a hammer and the resulting acceleration of the structure sensed with an accelerometer. The frequency components in the spectrum indicate mechanical resonances.

[6]More technically, the shape of each bandpass fi lter is the frequency response of the window function, translated to the centre frequency of that analysis band.

in audio signals tends to decrease at higher frequencies. A signal that sounds to the human ear as if it has flat frequency response will read as equal energy in constant percentage bandwidth filters.

To obtain this result with a digital computer, there are two possible approaches. One method is to create a constant-percentage bandwidth filter bank using digital signal processing methods which implement so-called FIR or IIR digital filters.

Another method is to use fast convolution. The signal is digitized and converted into spectral form using the FFT. The spectrum is then multiplied by functions that represent the various filter shapes and the energy determined in each filter bandwidth.

References

[1] *Vibration Spectrum Analysis: A Practical Approach*
Industrial Press, New York, 1991
Practical applications of spectral analysis in mechanical engineering.

[2] *Audio Spectrum Analyser*
Peter Hiscocks, Wireless World Magazine, May 1980, pp55-60
An audio-spectrum analyser that is suitable for construction by hobbyists.

[3] *Signal Processing Using Analog and Digital Techniques*
K.G.Beauchamp
George Allen & Unwin Ltd, 1973
Shows analog circuitry for various forms of spectrum analysis, which is a useful aid to understanding digital techniques.

[4] *Analog Integrated Circuit Design*
David A. Johns and Ken Martin
Wiley, 1997

[5] *Introductory Digital Signal Processing with Computer Applications*
Paul A. Lynn and Wolfgang Fuerst
Wiley, 1994
A very accessible introduction to digital signal processing, including digital spectrum analysis.

[6] `http://en.wikipedia.org/wiki/Window_function#Triangular_.28aka_Bartlett.29_window`
A collection of window functions.

[7] Stephan M. Bernsee
`http://www.dspdimension.com/data/index.html`
A very clear tutorial on the discrete fourier transform.