application
note

Introduction to FFT Analysis

By T Carey, Product Manager

This application note is an introduction to the basic principles of digital signal processing and gives some insight into how this technology is utilized within the IFR 2309 100 MHz to 2.4 GHz FFT Analyzer

Introduction

Digital Signal Processing (DSP) techniques when used in measurement instruments offer many advantages over traditional analog processing methods. DSPs are widely replacing analog systems due to their reliability, repeatability and programmability. As DSP devices have become more powerful and less expensive, their use in test instruments has become more widespread. 2309 from IFR is one such instrument which uses the Fast Fourier Transform (FFT) to translate digitized time domain signals into the frequency domain for the purpose of measurement. This application note introduces some of the basic concepts involved in digital signal processing such as analog to digital conversion and FFT, and then goes on to explain their use in IFR's 2309, 100 MHz to 2.4 GHz FFT Analyzer.

Analog to Digital Conversion Principles

Before applying DSP techniques to an analog signal, it must first be converted into a digital form. In general, this is accomplished using the process shown in figure 1. The aim is to convert a signal which is continuous in time and amplitude into one that is discrete in time and amplitude.

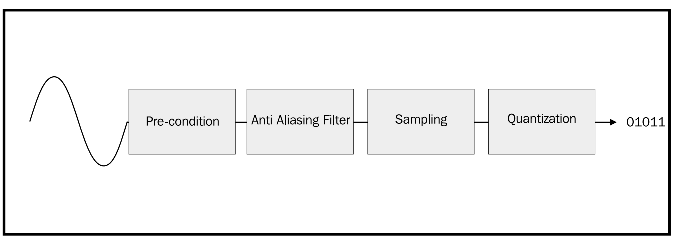

Figure 1 - Analog to Digital Conversion process

Sampling

In order to convert an analog signal into a digital one, it must first be sampled. This results are in natural sample values which remain continuous in amplitude but discrete in time.

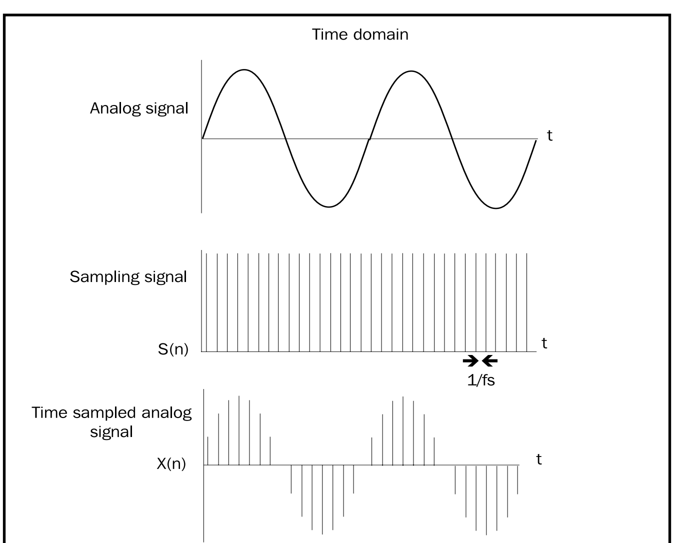

Figure 2 - Analog signal and sampled analog signal

Figure 2 shows a sinusoid both pre and post sampling. The sample values are exactly equal to the original signal value at the sampling instant. The amplitude of the analog signal is measured at discrete points in time using a sample and hold circuit.

The process of sampling leads to the frequency spectra of the baseband signal being reproduced around the sampling frequency.

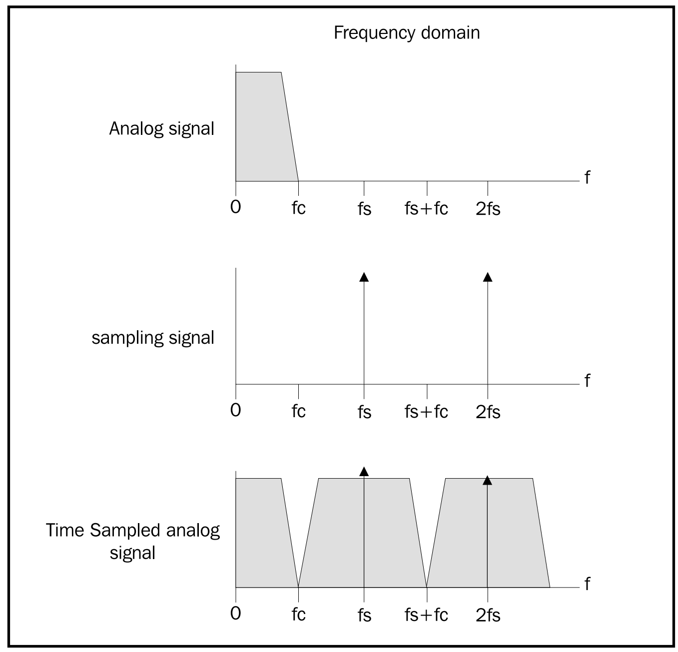

Figure 3a - Frequency domain of a sampled signal

Sampling Frequency

The choice of sampling frequency is very important. According to the sampling theorem, to be able to reconstruct a sampled signal accurately, the sampling frequency must be greater than twice the highest frequency contained in the signal to be sampled. This is called the Nyquist Frequency.

Nyquist frequency fn = 2* fc

where fc is the highest frequency present within the signal

Bandpass Sampling (IF sampling)

Instead of digitising a baseband signal it is often preferable to digitize an IF signal.. As an example consider a signal with a 300 kHz bandwidth at an IF of 10.7 MHz. According to the sampling theorem it would seem necessary to sample this signal at a rate at least 2 times the IF frequency i.e. 21.4 MHz. This would require a much more complex ADC. However, a closer look at the sampling theorem tells us that the sample rate need only be 2 times the bandwidth of the IF signal i.e. 600 kHz and hence the term IF or bandpass sampling.

The figure below shows the case where the signal to be sampled is at an IF rather than at baseband. In this case the sampling rate must be greater than 2 * (F_u - F_l).

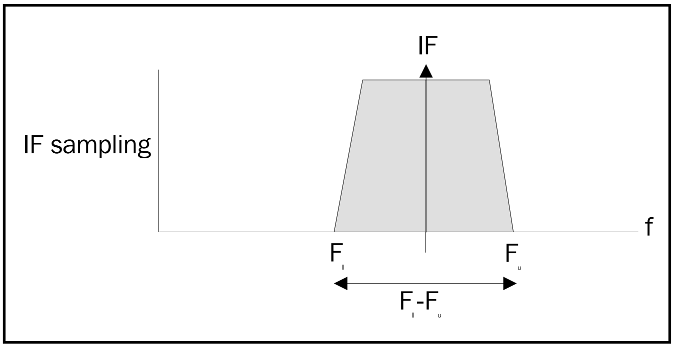

Figure 3b - IF Sampling or Bandpass sampling

Aliasing

If the analog signal is sampled at a frequency below fn, the Nyquist frequency, the samples taken from the signal can be used to construct a lower frequency 'alias' signal, as shown in figure 4 below. Here we have a 100 Hz sine wave sampled at 80 Hz i.e. very much less than 2 * fmax and as a consequence this signal could also be reconstructed as 20 Hz.

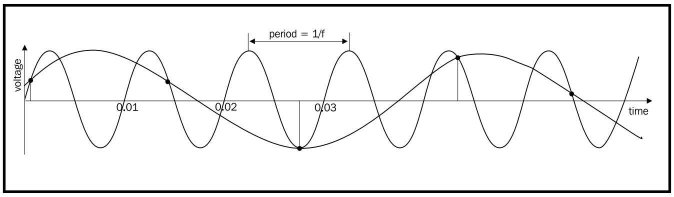

Figure 4 - A 100 Hz sine wave sampled at 80 Hz and a 20 Hz alias frequency

This effect is known as aliasing, and is combated by first ensuring the sampling rate is high enough and secondly through the use of an analog low pass filter in front of the ADC. This filter is called an anti-aliasing filter. Anti-aliasing filters attenuate frequency components above half the sampling frequency in the input signal to a level below the dynamic range of the analogue to digital converter (ADC). This ensures that any frequency components which could cause aliasing are suppressed. Anti-aliasing filters can theoretically be omitted if there is no possibility that frequency components above half the sampling frequency will be present. In practice this can rarely be guaranteed.

Signal Conditioning

It is important to ensure that the full voltage range of the ADC is used, so before a signal is input to the ADC, it is conditioned or scaled. Without pre-scaling, large signals can overdrive the converter leading to clipping and distortion. Conversely, it is important not to under drive the converter. Assuming an ADC has a full scale range of +/- 10 V, then without pre-scaling, a small 1 V input signal would not be digitally converted with the available resolution effectively utilising fewer bits than available and hence would suffer a poorer signal to noise ratio

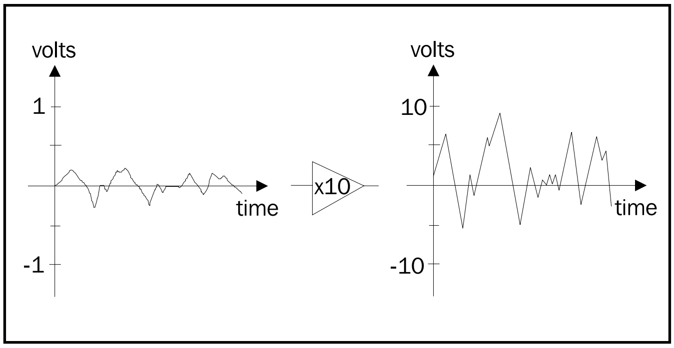

Figure 5 - Signal Conditioning

It is equally important to ensure that the voltage range of the input signal is not greater than that of the ADC, or clipping will occur. Any signal which falls outside the upper and lower voltage limits of the ADC will be converted as either the maximum or minimum voltage, so information is lost.

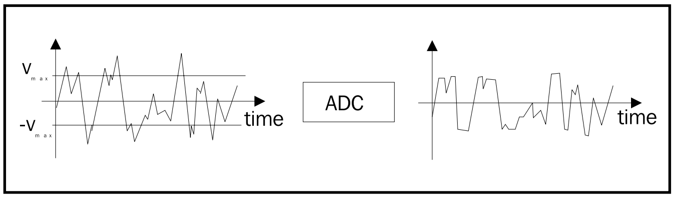

Figure 6 - Clipping

Quantization and Resolution

The output of the ADC consists of a fixed number of bits. The analog samples are quantized and mapped into their nearest digital value. The number of possible quantization levels is determined by the number of bits in the output of the ADC. For example, a 4 bit ADC can represent an analog sample as one of 16 possible values i.e. ($2^4 = 16$).

If the range of this converter is 0 and 10 volts, then its resolution is $10/(2^4-1) = 0.667$ volts/bit. The resolution of the converter is improved by increasing the number of bits while maintaining the voltage range of the converter. If the converter in the previous example were now changed to 8 bits, its resolution becomes $10/(2^8-1) = 0.039$ volts/bit, giving greater conversion accuracy.

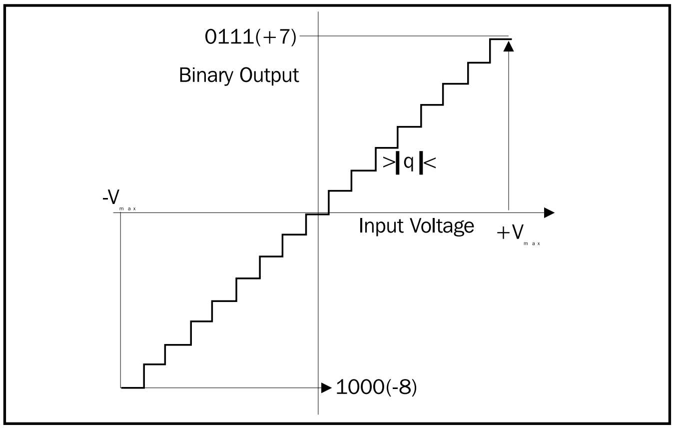

Figure 7 - 4 Bit Quantization

The time discrete natural sample values obtained earlier are now mapped to their nearest quantization level. This process introduces a non reversible error.

Quantization Error

If the value of the analog signal at a sample point falls between two quantization levels, the ADC treats it as if it were exactly one of the two values, thus introducing an error, known as quantization error. The maximum value of this error is half the value of the Least Significant Bit (LSB) of the ADC. By modelling the output of the ADC as a perfectly sampled signal and a noise signal (the quantization error), the signal to noise ratio of a multi-bit ADC can be calculated.

For real life signals, the quantization error signal is referred to as quantization noise. This noise signal has a rectangular probability distribution function of ± 1/2 LSB. This noise is a function of the input and can therefore be considered a distortion. The amount of distortion is given by D = q2/12 and is therefore governed by the number of quantization steps.

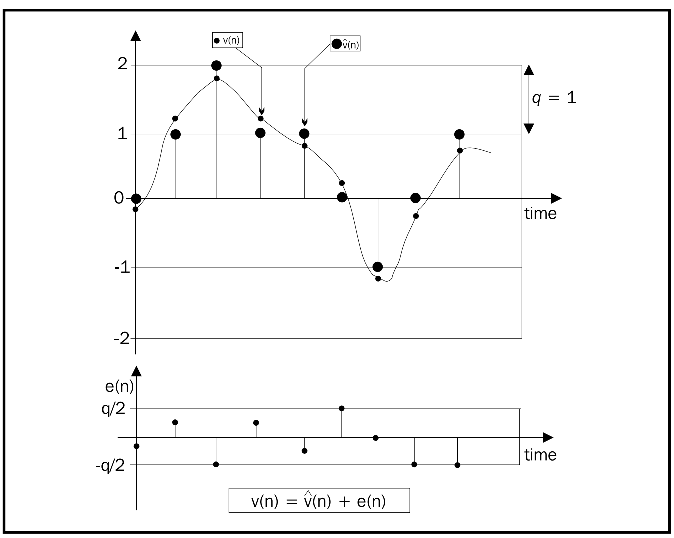

Figure 8 - ADC Output as perfectly sampled signal plus noise

Quantization noise is reduced by increasing the number of bits in the output of the ADC, which improves the resolution of the ADC. The '6 dB rule' for SNR in multi-bit ADCs illustrates this fact; SNR in dB = (n x 6.02)+ 1.76, where n is the number of bits used by the converter. An 8-bit ADC will therefore have an SNR of 49.92 dB. However, increasing the number of bits used by the converter also increases its complexity and cost.

Oversampling

The quantization noise is spread evenly across the range 0 to fs/2. If the sampling frequency is increased, the quantization noise must be spread across a wider frequency range and hence reduces it's level at every point. If a digital low pass filter is applied to the signal that passes the required baseband signal then the quantization noise falling outside the filter passband is rejected.

A reduction in quantization noise through oversampling and filtering has the same effect as using a ADC device with more quantization levels or bits. For each four fold increase in sampling rate above Nyquist, the effective number of bits, (ENOB) increases by 1 and so the dynamic range is improved by 6 dB.

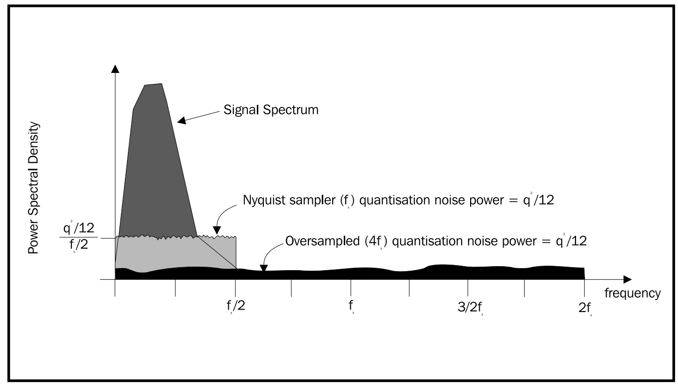

Figure 9 - The effect oversampling has on quantization noise

This approach can be used to reduce the number of bits required for a target dynamic range. In the limit, if the sampling rate is made very high the number of bits required for a given dynamic range requirement can be reduced to 1. This is the technique used in audio CDs and in IFR 2309 and is based upon Sigma Delta ADC technology.

Quantization Limits

Quantization can be a problem when sampling small-amplitude signals. If the signal is of the same order of magnitude as the quantization interval, the quantization process will produce a square wave output, or worse, a DC signal, as shown in figure 10.

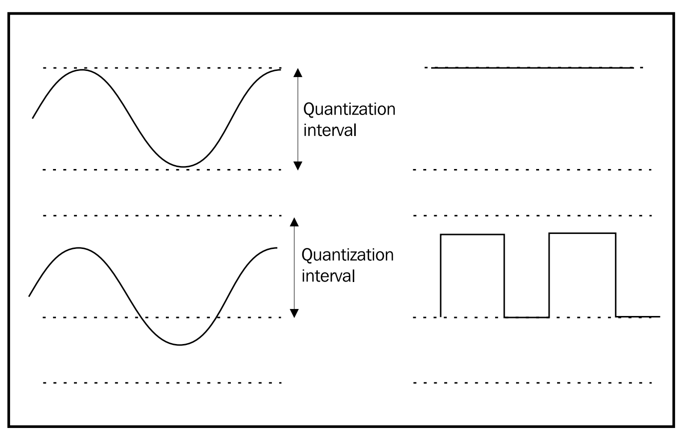

Figure 10 - Problems with quantization of small signals

One solution to this problem is to increase the resolution of the ADC by adding more quantization levels, i.e. more bits. However, this is often not practical, either because of the high signal frequency or for reasons of cost, so a technique known as dithering is used. This involves adding noise to the input signal before quantization. The result of this is a Pulse Width

Modulated (PWM) signal which can be decoded by averaging to give the original signal.

Dithering

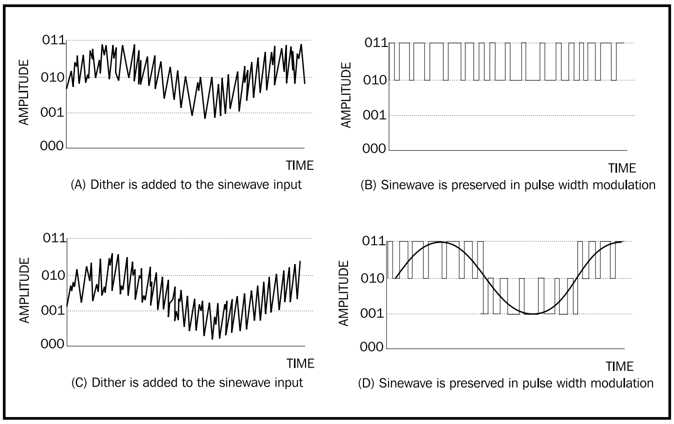

Figure 11 - Small signal quantization with dithering

Dithering has the effect of making the ADC more linear at low input levels. However, this increase in linearity comes with an increase in the level of noise in the ADC. Care must be taken when choosing the level of dithering to ensure that the dynamic range of the ADC is not compromized.

IFR 2309

2309 takes a different approach to the process of analog to digital conversion, employing a patented single-bit bandpass sigma delta converter. This converter resolves the problems inherent in using multi-bit ADCs for high-speed conversion, while offering excellent linearity and improved accuracy over conventional ADC methods. (refer to the sigma delta converter application note).

2309 samples the signal at a rate of 40.32 MHz to reduce quantization noise. A 300 kHz bandwidth is sampled around the 10.71 MHz IF. This provides the best compromize between resolution bandwidth and signal to noise ratio in the instrument.

An anti-aliasing filter with very high rejection (100 dB) complements the wide dynamic range provided by the converter.

The single bit converter achieves very high linearity leading to accurate ratiometric measurement such as required when measuring intermodulation and spurious signals in amplifiers and frequency converters.

THE FAST FOURIER TRANSFORM

Introduction to Fourier Transforms

Fourier Series

Any periodic waveform can be decomposed into a series of sine and cosine waves. The general expression is shown below.

$$g(t) = \sum_{n=0}^{\infty} A_n \cos\left(\frac{2\pi nt}{T}\right) + \sum_{n=0}^{\infty} B_n \sin\left(\frac{2\pi nt}{T}\right)$$

The Fourier Transform (FT) is a mathematical operation which converts time-domain signals to the frequency domain. Being able to visualize a signal in the frequency domain offers many benefits over time-domain representations. Frequency domain representations allows individual frequency components contained within a signal to be viewed including modulation sidebands, distortion effects and spurious frequency components.

A good example of this is the square wave of figure 12. A square wave is composed of a fundamental frequency and all its odd harmonics. The level of each harmonic decreases in amplitude with harmonic number. Viewed in the time domain representation, the square wave is gives no indication of its composition. Viewed in the frequency domain all frequencies components are displayed along with their relative amplitude.

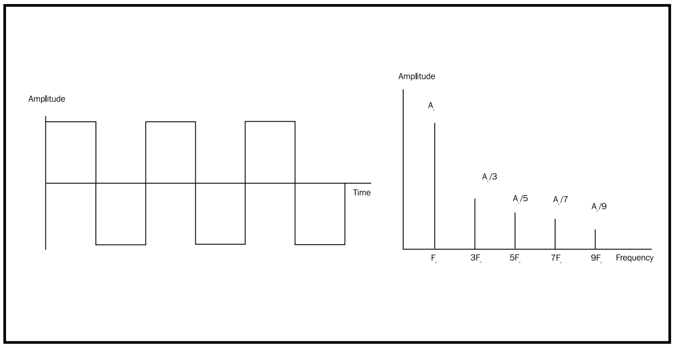

Figure 12 - Time- and frequency-domain representations of a square-wave signal

In addition, the FT conserves the phase information contained in the signal, which can be used to measure the relative phase of different frequency components, or the phase difference between a number of signals acquired simultaneously.

The Discrete Fourier Transform (DFT) is a version of the FT which can be applied to sampled time-domain signals. The DFT produces a discrete frequency spectrum; i.e. amplitude levels at discrete frequencies, or 'frequency bins'. The DFT is defined as:

$$x(f) = \sum_{n=0}^{N} x(n) e^{\frac{-j2\pi fn}{f_s}}$$

N is the total number of samples taken from the original signal. Essentially this equation uses the values, x, of the N samples to calculate the amplitude, X, of the signal at the kth discrete frequency bin. The total frequency spectrum is constructed from all these values over the whole frequency range.

Instruments use digital signal processing (DSP) hardware to apply this algorithm to convert the sampled time-domain signal into a frequency-domain one. However, when there are a large number of samples, as is often the case in order to obtain a good representation of the signal, the speed of the conversion process suffers.

Fast Fourier Transform

The Fast Fourier Transform (FFT) is a development of the DFT which removes duplicated terms in the mathematical algorithm to reduce the number of mathematical operations performed. In this way, it is possible to use large numbers of samples without compromising the speed of the transformation. The FFT reduces computation by a factor of $N/\log_2 N$. The relationship between frequency span, sampling frequency, resolution and display time in a multi-bit FFT analyzer, for a typical sampling frequency of 2.56 times the signal frequency, is summarized in table 1, for FFTs of different lengths N.

Frequency Span (f)	***Sampling Frequency (f_s)***	***Frequency Resolution ($R= f_s/N$)***	***Update Time ($T = N/f_s$)***
1024 Point FFT			
1 Hz	2.56 Hz	2.5 mHz	400 s
10 Hz	25.6 Hz	25 mHz	40 s
100 Hz	256 Hz	0.25 Hz	4 s
1 kHz	2.56 kHz	2.5 Hz	0.4 s
10 kHz	25.6 kHz	25 Hz	40 ms
100 kHz	256 kHz	250 Hz	4 ms
300 kHz	768 kHz	750 Hz	1.33 ms
2048 Point FFT			
1 Hz	2.56 Hz	1.25 mHz	800 s
10 Hz	25.6 Hz	12.5 mHz	80 s
100 Hz	256 Hz	0.125 Hz	8 s
1 kHz	2.56 kHz	1.25 Hz	0.8 s
10 kHz	25.6 kHz	12.5 Hz	80 ms
100 kHz	256 kHz	125 Hz	8 ms
300 kHz	768 kHz	375 Hz	2.67 ms

Table 1 - Sampling frequency, display update time and resolution for different frequency spans with differing FFT lengths.

2309 is able to display frequency spans between 10 Hz and 300 kHz, using a 2048 point FFT. A dedicated Harris decimator chip is used to reduce the 40.32 MHz sampled signal to 630 kHz for 2309's maximum frequency span of 300 kHz. Further hardware decimation is used for spans down to 39.375 kHz. In fact, decimation is a slightly misleading term, since the frequency is reduced by powers of two by only taking every second, fourth, eighth etc sample. Below this, software routines are used to reduce the sampling rate. By using this oversampling and decimation, quantization noise is reduced, and accuracy and resolution are improved.

Frequency Span	***Sampling Rate***	***Frequency Resolution***	***Update Time***
300 - 157.5 kHz	630 kHz	307.6 Hz	~100 ms
157.5 - 78.75 kHz	315 kHz	153.8 Hz	
78.75 - 39.375 kHz	157.5 kHz	76.9 Hz	

Table 2 - Sampling frequency, display update time and resolution for different frequency spans on 2309.

Spectral Leakage

This is a phenomenon which occurs because the FFT algorithm can only be applied to periodic signals so the sampled input signal is 'periodized'. If, as in most cases, the sampled signal is not periodic, or an integer number of periods is not sampled, discontinuities occur in the periodic signal processed by the FFT, causing the energy contained in the signal to 'leak' from the signal frequency bin into adjacent frequency bins. This leakage causes amplitude errors in the frequency spectrum display. Figure 13 shows a periodized time domain signal where an integer number of cycles has been sampled and one where a non-integer number of cycles has been sampled. Frequency spectra for integer and non-integer sampled signals are shown in figure 14.

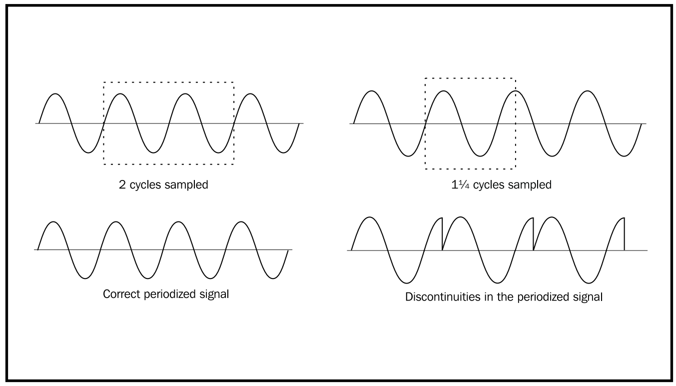

Figure 13 - Periodized time domain signal produced from integer and non-integer samples

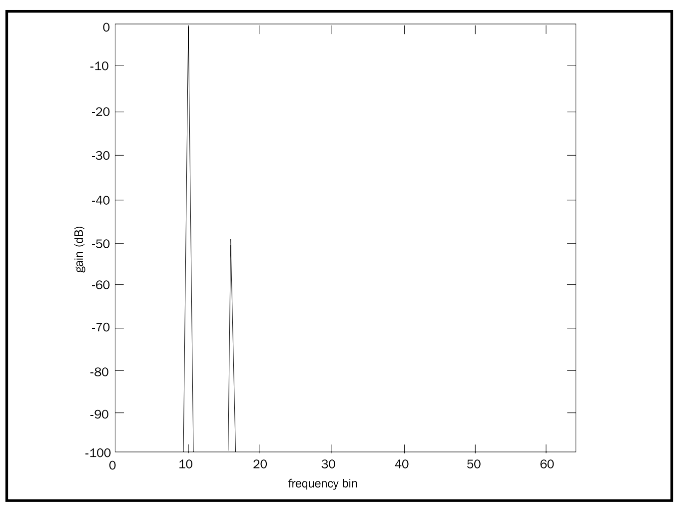

Figure 14 - Spectrum display without spectral leakage

As a result of the amplitude errors caused by spectral leakage, small frequency peaks occurring close to larger ones may be obscured, as shown in figure 15.

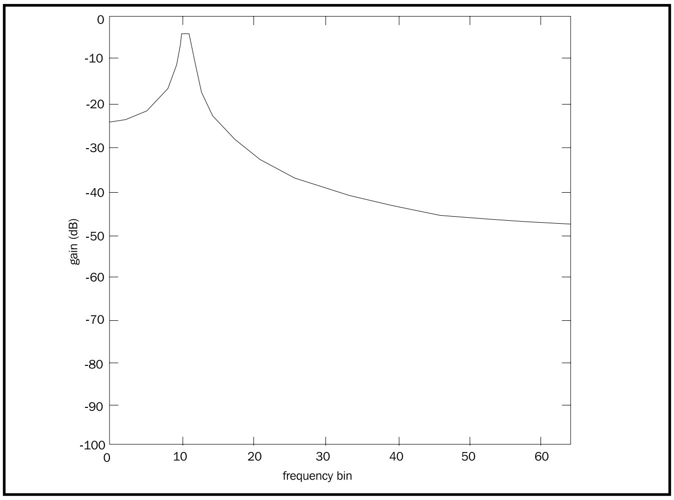

Figure 15 - Small frequency peak obscured as a result of leakage

Windowing

To reduce the effects of spectral leakage, window functions are used. The effect of a window is to assign a weighting coefficient to each of the input samples, reducing those samples that cause spectral leakage. In effect, samples at the beginning and end of the sampling period are reduced to zero so that the discontinuities in the periodized sampled signal are removed. The effect of windowing on a signal is shown in figure 16.

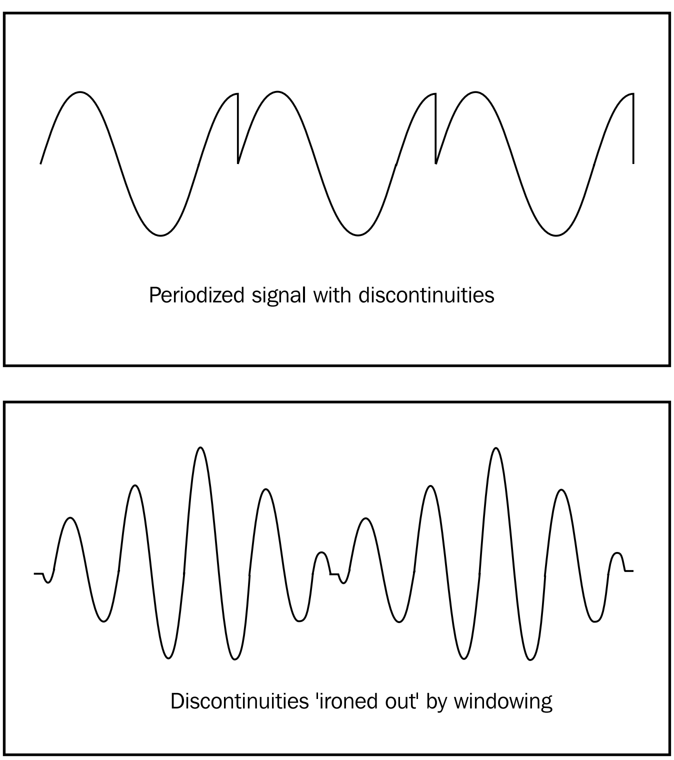

Figure 16 - The effect of windowing on a time-domain signal

Figure 17 shows the effect of windowing on the frequency-domain representation of the two tone signal shown in figure 14. It can be seen that the smaller frequency peak is no longer obscured.

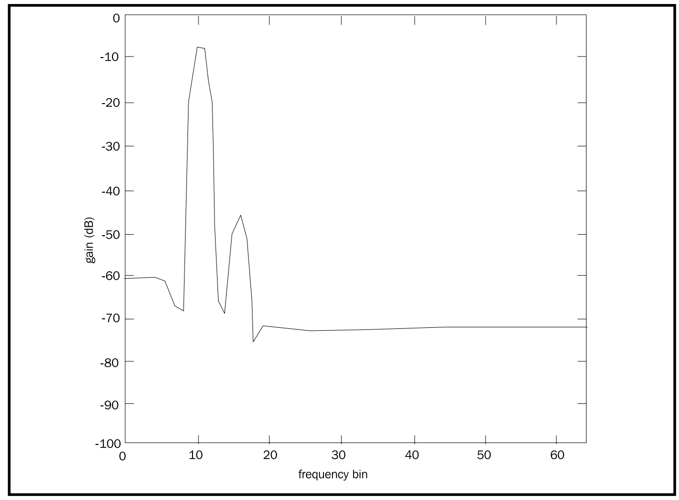

Figure 17 - Adjacent peak spectrum display with windowing

Windows are characterized by a number of properties, as shown in figure 18.

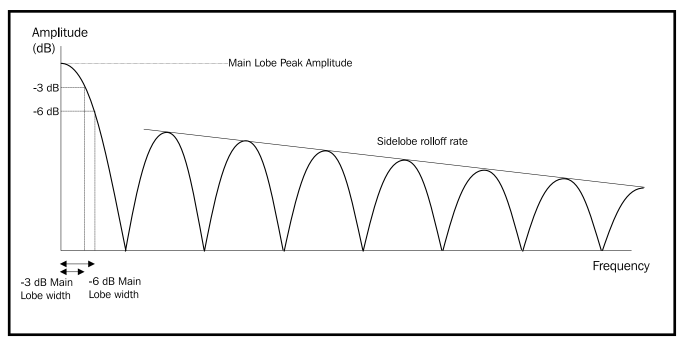

Figure 18 - Spectrum of a window

The shape of the window's main lobe is defined by the -3 dB and -6 dB main lobe width. These are defined as the width of the main lobe, in frequency bins, where the window response becomes respectively 3 dB or 6 dB less than the main lobe peak gain. The width of the main lobe in the frequency spectrum is important, as it affects the frequency resolution of the window (i.e. ability to distinguish between closely spaced frequency components). As the main lobe narrows, frequency resolution increases. However, with this narrowing of the main lobe, the window energy spreads into the side lobes, increasing the spectral leakage. Therefore, a compromize between frequency resolution and spectral leakage must be reached.

Maximum side lobe level is defined as the level, in decibels, of the maximum side lobe, relative to the main lobe peak gain.

Side lobe rolloff rate is the rate of decay of frequency of the side lobe peaks, in decibels per decade. Table 2 lists the characteristics of some common window functions.

Window	-3 dB Main Lobe Width (bins)	-6 dB Main Lobe Width (bins)	Maximum Side Lobe Level (dB)	Side Lobe Rolloff Rate (dB/dec)
Uniform (None)	0.88	1.21	-13	20
Hanning	1.44	2.00	-32	60
Hamming	1.30	1.81	-43	20
Blackman-Harris	1.66	1.81	-71	20
Exact Blackman	1.52	2.13	-67	20
Blackman	1.68	2.35	-58	60
Gaussian		2.18	-55	20
Flat Top	2.94	3.56	-44	20

Table 3 - Characteristics of common window functions

The choice of window depends upon the frequency content of the signal. A popular choice is the Hanning window. This window has quite a narrow main lobe, therefore good frequency resolution, and reasonable side lobe suppression, making it suitable for many applications. 2309 offers a choice of two windows: a 5 term Blackman-Harris designed window, giving excellent sideband rejection with an acceptably narrow main lobe, or a Gaussian window, with a slightly wider main lobe and less sideband rejection.

Scallop Loss

2309 also employs a unique patented algorithm, Gaussian Term Interpolation (GTi) to reduce amplitude errors caused by scallop loss. Scallop loss is an error caused by the discrete nature of the frequency spectrum; the signal is displayed as amplitude levels at equally spaced discrete frequency 'bins'. If the signal frequency coincides with the center of one of the discrete frequency bins, the correct peak level is displayed. If, however, the signal frequency is not at the center of a frequency bin, then a reduced level is displayed, causing an error which can be up to 3 dB. Flat top windows have traditionally been used to reduce scallop loss, but these have a wide main lobe, which compromizes resolution, and poor sidelobe rejection. The GTi algorithm employed in 2309 uses the incorrect peak level and the level in the two adjacent frequency bins to estimate the correction required to the peak level, and adjusts the peak level accordingly. In this way, 2309's frequency resolution and sidelobe rejection are not compromised.

2309

2309 from IFR is an RF FFT analyzer operating between 100 MHz and 2.4 GHz offering numerous benefits over traditional swept tuned spectrum analyzers, particularly regarding measurement speed, dynamic range, accuracy and linearity.

2309 is based on the architecture shown in figure 19. By using FFT and DSP techniques, 2309 is able to offer improved speed and accuracy over conventional frequency analysis instruments. The unique patented sigma delta ADC, coupled with high quality downconversion and low noise RF synthesis give 2309 its wide dynamic range and exceptional linearity.

Conclusion

The basics of many of the methods used in digital signal processing and frequency analysis have been introduced, and their use in 2309 explained.

Thanks to its unique architecture, 2309 is able to offer many benefits over conventional swept tuned spectrum analyzers. Its increased speed, accuracy, dynamic range and detector linearity offer new possibilities for measurement.

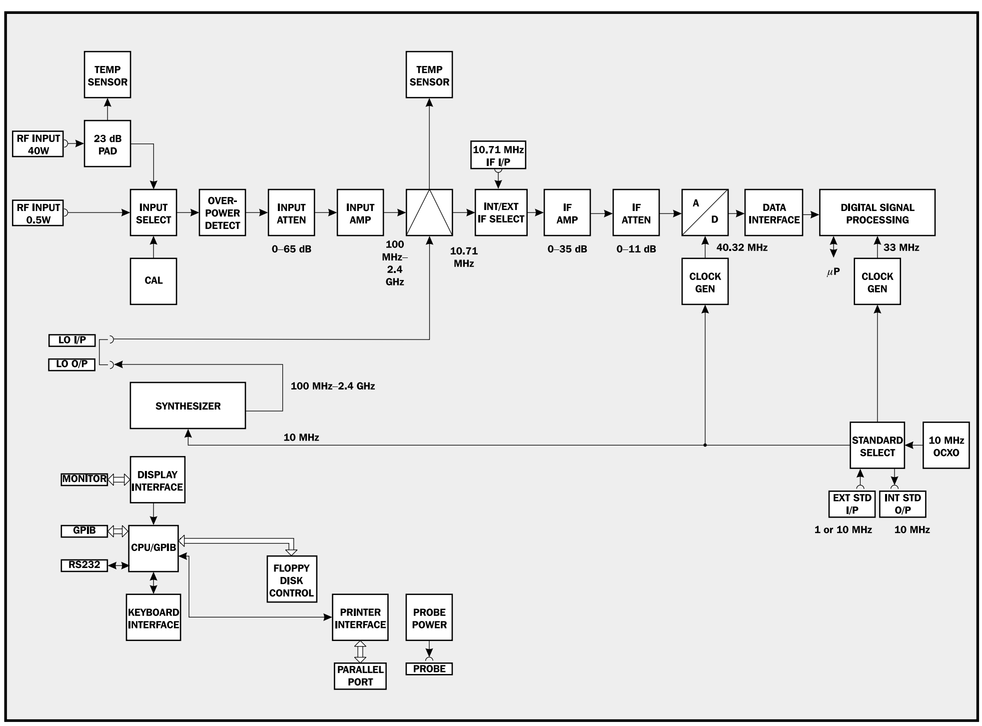

Figure 19 - 2309 Block diagram

IFR Americas, Inc., 10200 West York Street, Wichita, Kansas, 67215-8999, USA. E-mail: info@ifrsys.com
Tel: +1 316 522 4981 Toll Free USA: 1 800 835 2352 Fax: +1 316 522 1360

IFR, Longacres House, Six Hills Way, Stevenage SG1 2AN, United Kingdom. E-mail: info@ifrsys.com
Tel: +44 (0) 1438 742200 Freephone UK: 0800 282 388 Fax: +44 (0) 1438 727601

As we are always seeking to improve our products, the information in this document gives only a general indication of the product capacity, performance and suitability, none of which shall form part of any contract. We reserve the right to make design changes without notice. All trademarks are acknowledged. Parent Company IFR Systems, Inc. © IFR Ltd. 2001.