TECHNICAL REVIEW

No. 1 — 1979

Contents

The Discrete Fourier Transform and FFT Analysers

by

N. Thrane, Ph. D.

ABSTRACT

The FFT approach to frequency analysis is in principle completely different from that of analog or digital filters. In FFT analysers *multiplication* is performed between the input signal and a weighting function in the time domain, while a filtration of the signal involves *convolution* between the input signal and the impulse response of the filter in the time domain.
The article discusses in detail the assumptions on which the FFT analysis is based and their effects, however in a very pictorial and non-mathematical way. It also shows how these effects can be avoided or diminished by use of anti-aliasing filters and proper selection of the time weighting function.

SOMMAIRE

La transformation de Fourier rapide (FFT) est une approche de l'analyse de fréquence totalement différente du filtrage numérique. Dans les analyseurs FFT on effectue une *multiplication* dans le domaine temps du signal d'entrée par une fonction de pondération, alors que le filtrage du signal demande une *convolution* dans le domaine temps du signal d'entrée par la réponse impulsionnelle du filtre.
Cet article discute en détails des hypothèses sur lesquelles est basée l'analyse FFT et leurs effets, mais d'une façon imagée et non pas mathématique. On montre également comment ces effets peuvent être éliminés ou diminués en utilisant des filtres anti-repliement et en choisissant de façon adéquate la fonction de pondération temporelle.

ZUSAMMENFASSUNG

Die schnelle Fouriertransformation (FFT) bewältigt die Frequenzanalyse nach einem vollkommen anderen Prinzip als analoge oder digitale Filter. Während FFT-Analysatoren eine *Multiplikation* in der Zeitdomäne zwischen dem Eingangssignal und einer Zeitbewertungsfunktion durchführen, führt das Filterprinzip eine *Faltung* in der Zeitdomäne zwischen dem Eingangssignal und der Impulsantwort des Filters mit sich.

In diesem Artikel werden an Hand von vielen Bildern und mit wenig Mathematik die Voraussetzungen auf denen FFT basiert, sowie auch ihre Effekte detailiert diskutiert. Außerdem wird gezeigt, wie sich diese Effekte durch Verwendung von Antialiasing-Filtern und Wahl geeigneter Zeitbewertungsfunktionen vermieden bzw. unterdrücken lassen.

Introduction

Modern digital FFT analysers work on the principle of a fast and efficient calculation of the so-called Discrete Fourier Transform. This special version of the Fourier integrals is discrete, i.e. sampled, in both time and frequency domain and hence lends itself to a direct calculation in a digital processor.

It is important to realise that the Discrete transform handles the time signal in a way quite different from that found in more traditional analysis. When a frequency analysis is performed by use of filters, analog or digital, there is a steady flow of time data into the filter, which in turn produces a steady flow of filtered time data at the output. In this way the time signal is continuously being processed.

The Discrete transform, on the other hand, processes the time signal in blocks of data. Samples of the time signal are stored in a digital memory, and when this is filled up, the whole memory is transformed into the frequency domain as one block. Furthermore, in order to obtain discrete frequency components, it is assumed that this block represents one period of a periodic signal. Actually, this is very similar to the way a Time Compression Analyser works. There a digital memory is repetitively played back, producing a periodic signal, which is then analysed by use of a sweeping filter.

It is therefore a common property of Time Compression and FFT analysers that they work only on a small part of the time signal, and treat this as part of a periodic signal. Hence, the original input signal has been "time limited" before the analysis. This time limitation has important consequences on the result of the analysis and may lead to some unexpected, and even peculiar, effects. Although these effects can also be observed with time compression analysers, they are even more pronounced with an FFT analyser. In the following discussions we shall therefore concentrate on the way this time limitation is performed, i.e. the use of different time weighting functions, and also explain the various effects graphically instead of mathematically.

The following discussions will be very brief and qualitative. For a more quantitative discussion of the Fourier Transform, and especially subjects like complex notation, sampling and convolution theorems, the reader is referred to references [1] to [3] given at the end of the article. A general discussion of the advantages of digital analysis over analog analysis can be found in reference [4] and therefore will not be dealt with here.

The Fourier Transform

Before discussing the Discrete Fourier Transform in detail it will be useful to discuss the various transforms shown in Fig.1.

Firstly, in Fig.1a we have the Integral transform, which transforms a continuous time signal extending over all time, $-\infty < t < +\infty$, into a continuous frequency spectrum also extending over all frequencies, $-\infty < f < +\infty$. It can be said that this is the ideal transform, which in principle could be applied to all practical signals. However, since it requires knowledge of all of the time signal it can in practice only be applied to relatively short transient signals while "continuous" signals must be handled by other means.

Secondly, in Fig.1b we have the well-known Fourier Series, which apply to periodic time signals. Hence, only one period of the time signal has to be specified and included in the transform. In this case we find that a periodic and continuous time signal is transformed into a discrete frequency spectrum, exhibiting all the harmonic frequencies.

Thirdly, in Fig.1c we find the transform which applies to sampled time signals. This is actually exactly the opposite situation of Fig.1b. Here the discrete samples of the time domain are transformed into a periodicity of the frequency spectrum, thus demonstrating the basic symmetry of the Fourier transform between time and frequency. A periodicity in one domain corresponds to discrete samples in the other domain. It is of interest to note how the sampling of the time signal leads to aliasing of frequencies. Due to the symmetry and periodicity of the frequency spectrum a component of frequency f_c in the continuous signal will appear in the sampled signal at frequencies $f_d = nf_s \pm f_c$, where f_s is the sampling frequency and $n = 0, \pm 1, \pm 2.....$ In order to avoid ambiguities about the frequency content of the continuous signal it is therefore necessary to band-limit the input signal to frequencies less than half the sampling frequency. This is actually the content of the sampling theorem. If the input signal itself is not band-limited as prescribed, it should be passed through a low-pass antialiasing filter before being sampled.

1. Integral Transform (a)

$$G(f) = \int_{-\infty}^{\infty} g(t) e^{-j\,2\pi ft}\, dt$$

$$g(t) = \int_{-\infty}^{\infty} G(f)\, e^{j2\pi ft}\, df$$

Infinite and continuous in time and frequency domains

2. Fourier Series (b)

$$G(f_k) = \frac{1}{T}\int_{-\frac{T}{2}}^{\frac{T}{2}} g(t) e^{-j2\pi f_k t}\, dt$$

$$g(t) = \sum_{k=-\infty}^{\infty} G(f_k) e^{j2\pi f_k t}$$

Periodic in time domain
Discrete in frequency domain

3. Sampled functions (c)

$$G(f) = \sum_{n=-\infty}^{\infty} g(t_n) e^{-j2\pi f t_n}$$

$$g(t_n) = \frac{1}{f_s}\int_{-f_s/2}^{f_s/2} G(f) e^{j2\pi f t_n}\, df$$

Discrete in time domain
Periodic in frequency domain

4. Discrete Fourier Transform (d)

$$G(f_k) = \frac{1}{N}\sum_{n=o}^{N-1} g(t_n) e^{-j\frac{2\pi nk}{N}}$$

$$g(t_n) = \sum_{k=o}^{N-1} G(f_k) e^{j\frac{2\pi nk}{N}}$$

Discrete and periodic in both time and frequency domains

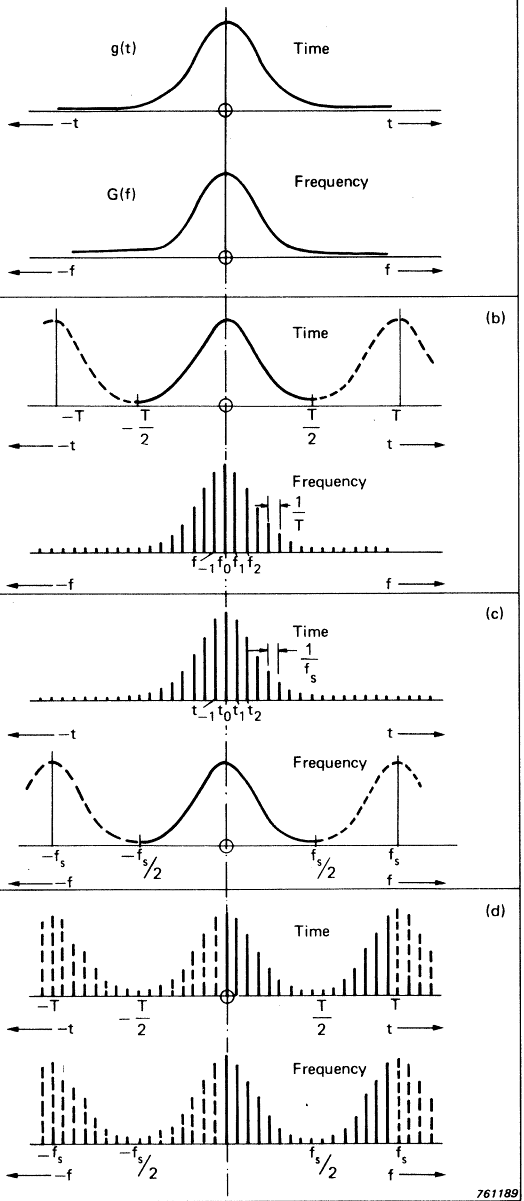

Fig. 1. Various forms of the Fourier Transform

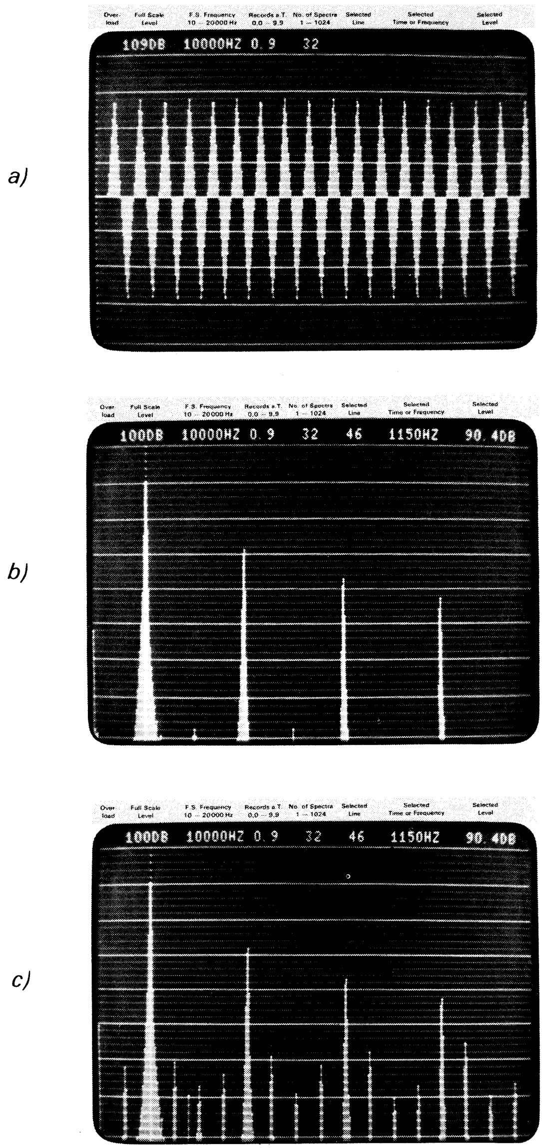

Fig. 2. FFT analysis a) Triangular time signal RMS power spectrum b) with antialiasing filter c) without antialiasing filter

In Fig.2 is shown the result of an FFT analysis* of a triangular time signal, which contains odd harmonics far above the sampling frequency. The frequency range shown corresponds to $\approx 0{,}4\,f_s$. The correct analysis is obtained when the signal has been passed through an antialiasing filter with a cut-off frequency of $0{,}4\,f_s$ prior to the sampling. The same signal was also analysed without using the antialiasing filter, clearly showing the higher harmonics appearing at their alias-frequencies falling within the analysis range.

The Discrete Fourier Transform

In Fig.1d we have finally the Discrete Fourier Transform, which applies to a discrete *and* periodic time signal. Accordingly, the frequency spectrum must also be periodic and discrete. Due to the periodicity in both time and frequency domain only a finite number of samples are involved in the transform, and therefore this particular transform can be calculated directly by digital means. Actually, if one period of the time signal is described by N samples, the frequency spectrum will also contain N samples per period. However, due to the symmetry of the frequency spectrum (which is the case for all real time signals) only N/2 of the frequency samples will be independent. In addition, the use of an antialiasing filter with a cut-off frequency below $f_s/2$ will reduce the number of valid frequency components further. Typically, a time signal having 1024 samples will allow for the calculation of a 400 line frequency spectrum.

To help better understanding of the Discrete transform it will be derived qualitatively and pictorially from the integral transform. This is shown in Fig.3, and involves three steps: 1) Time sampling, 2) Time limitation, and 3) Time convolution or frequency sampling. These steps are not performed in an FFT analyser but serve only to clarify the differences between the integral and the discrete transforms.

---

* The photos shown in Fig.2 and in the following illustrations were obtained directly from the display of an FFT analyser. The Brüel & Kjær Narrow Band Analyser Type 2031 was used. In all cases the frequency range was 0 — 10 kHz, corresponding to a sampling frequency of 25,6 kHz. The horizontal frequency axis is linear, while the vertical amplitude axis is logarithmic, covering a range of 80 dB. The frequency spectrum (RMS power spectrum) is shown as 400 lines with a separation of 25 Hz. The sampled time signals used for the analysis are also displayed. Here, both amplitude and time axis are linear. The 1024 time samples in this case represent 40 ms of the input signal. Note, that in Fig.2a only part of the time signal is displayed.

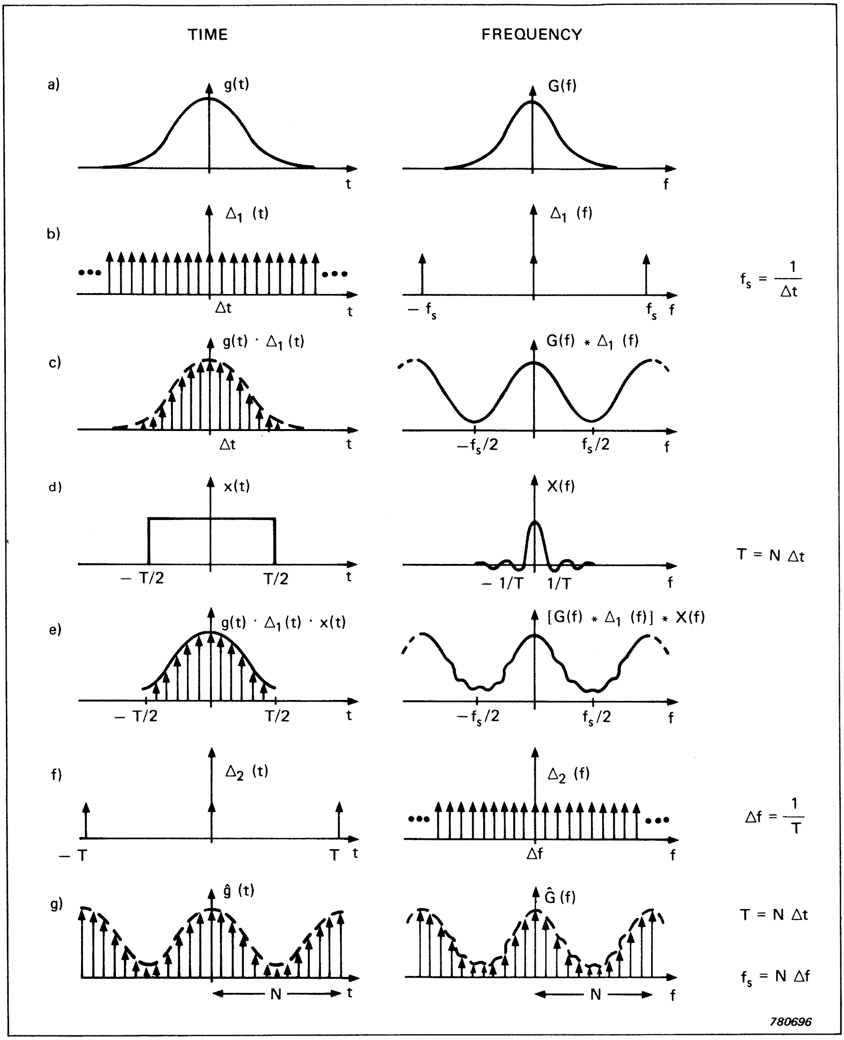

Fig.3. Derivation of the Discrete Fourier Transform from the Integral Transform

1) Time sampling

The continuous time signal and its frequency spectrum obtained using the Integral Transform are shown in Fig.3a. The time signal is sampled by multiplying it with the sampling function of Fig.3b. This function consists of an infinite series of impulses with a separation of Δt. The Fourier transform of the sampling function is another series of impulses with a separation of $f_s = 1/\Delta t$. According to the convolution theorem, multiplication of the two time signals will correspond to a convolution of their respective frequency spectra. The results are shown in Fig.3c, where the periodicity of the frequency spectrum corresponding to the sampled time signal can be recognised. Hence, the time sampling gives rise to the possibility of aliasing of frequencies. However, it can be avoided by the correct use of an antialiasing filter, as discussed above.

2) Time limitation

In order to limit the number of samples in the time function the sampled time signal (Fig.3c) is multiplied by a time weighting function or time window. This is shown in Fig.3d as a rectangular signal of length T. The Fourier transform of this time window is the well-known sinx/x function also shown. The result of the time window multiplication is given in Fig.3e, where the number of samples, N, in the time limited signal is given by $N = T/\Delta t$. In the frequency domain the time window multiplication is again reflected as a convolution of the two frequency spectra, introducing ripples into the frequency spectrum. If the original time signal had included discrete frequency components, these would now have been replaced by a sinx/x function. Hence, the sinx/x function actually determines the filter characteristic of the analysis, giving rise to a leakage of power from its original frequency into the neighbouring frequencies. This effect of the time limitation is often named ripples, sidelobes or leakage, and is basically determined by the Fourier transform of the time window used. In the case of a rectangular time window the filter characteristic (sinx/x) will have a very poor selectivity and limit the useful dynamic range to less than 40 dB. However, since the high levels of the sidelobes of the sinx/x are determined by the discontinuities of the rectangular time window, the use of a continuous and smooth time window, e.g. the Hanning window,* will diminish this effect. Since only samples of the filter characteristics are actually measured with an FFT analyser, a more detailed discussion of this effect will be dealt with later.

3) Time convolution

The time signal shown in Fig.3e can be used for a digital calculation.

* The Hanning window consists of one period of a cosine signal of length T, displaced so that it starts and stops at zero, see Fig.4.

However, the corresponding frequency spectrum, being continuous, cannot be calculated. What is required is a sampling of the frequency spectrum which can be achieved, as shown previously, by multiplying with a sampling function in the frequency domain. The sampling function is shown in Fig.3f where the impulses are separated by $\Delta f = 1/T$, corresponding to N samples within one period of the frequency spectrum ($f_S/\Delta f = f_S \times T = T/\Delta t = N$). In the time domain the frequency sampling corresponds to a convolution of the time limited signal with impulses of a separation of T. Hence, the effect of this is to produce a periodic time signal with the time limited part as one period. The final result — The Discrete Fourier Transform — is shown in Fig.3g.

Sampling in the frequency domain is associated with the so-called "Picket-fence effect". Since we do not take the full continuous spectrum into account, but only samples of it, this corresponds to observing the spectrum through a picket-fence. If there are very peaked components in the spectrum we might not observe the correct maximum value, but only the lower values on the slopes of the peak. This is by no means a new effect. It is found whenever discrete filters are used for frequency analysis, e.g. also with 1/3 octave filters. The amplitude errors which occur depend on the amount of overlap between adjacent filters. In the case of an FFT analyser using the rectangular time window the error will be less than 3,9 dB. However, using the smoother Hanning window the error is reduced to a maximum of 1,4 dB. Since the smooth Hanning window has an effective length shorter than the rectangular window, the bandwidth of its Fourier transform will be wider than that of the sinx/x function, and hence give a better overlap of the filter characteristics. Thus the Hanning window will compensate both for the sidelobe effect and for the picket-fence effect.

It should be noted that the Discrete transform produces frequency data corresponding to filters of constant absolute bandwidth equidistantly placed along a linear frequency axis. Furthermore, the data represents the maximum information which can be obtained from a time signal of length T, i.e. the BT product is unity.

When the discrete transform is used in an FFT analyser, the continuous time signal is sampled and time limited, as shown in Fig.3, using an appropriate time window. However, it is important to note that when the N samples are stored in the digital memory, no recirculation of this memory is performed in order to produce a periodic signal. It is an *assumption* of the discrete transform that the N samples used for the mathematical calculations describe *one* period of a fictitious periodic signal.

How well the frequency spectrum of this fictitious signal corresponds with that of the original continuous signal all depends on the approximations involved, i.e on the aliasing of frequencies, the sidelobe effect and the picket-fence effect.

Time Weighting Functions

The effect of time limiting the signal on its frequency spectrum will now be discussed, i.e. the use of the different time weighting functions. In this discussion only the two most important functions will be considered: the rectangular, or flat, window and the smooth Hanning window. Although other weighting functions exist we shall not lose in generality by this restriction.

In Fig.4 are shown the filter characteristics of the two window functions, i.e. the numerical value of their Fourier transforms. The main features of the filter characteristics are summarised in Table 1. It is clearly seen that the Hanning window gives a characteristic with much better selectivity than the flat window. Not only is the highest sidelobe nearly 20 dB lower in amplitude but also the rate of fall of the sidelobes is much higher. In practice the Hanning window will therefore permit the use of a much higher dynamic range. However, the price to be paid for these advantages is in terms of an increased bandwidth. As mentioned

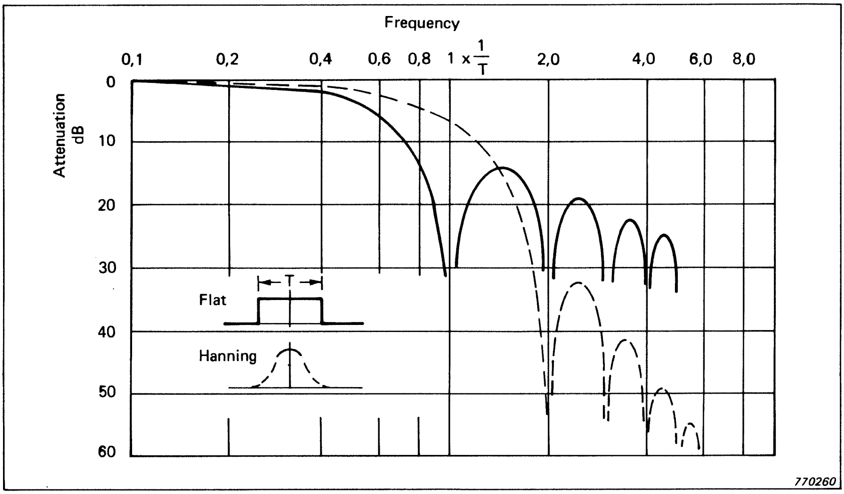

Fig.4. Comparison of the spectra (filter characteristics) of the rectangular and the Hanning time weighting functions

Name	Noise bandwidth	Highest sidelobe	Sidelobe fall-off rate
Flat Hanning	$1 \times \Delta f$ $1{,}5 \times \Delta f$	−13 dB −32 dB	20 dB/decade 60 dB/decade
$\Delta f = 1/T$ = line separation in sampled spectra			

790107

Table 1 Comparison of window functions

earlier, this actually reduces the picket-fence effect to a more acceptable level, and could therefore be regarded as an advantage too. When measuring power spectral densities it is necessary to divide the power spectrum with the filter bandwidth, and one will therefore find a difference of a factor of 1,5 (1,76 dB) between measurements performed with the flat and the Hanning window. The peak amplitude of the Hanning window is often chosen as 2. In this way sinusoidal signals will give the same peak amplitude independent of the weighting function.

It is interesting to study the effect of time limitation on the simple example of a sinusoidal signal. This is demonstrated in Fig.5. According to the Integral Fourier Transform the spectrum of such a tone burst consists of two sinx/x functions, situated at the positive and negative fre-

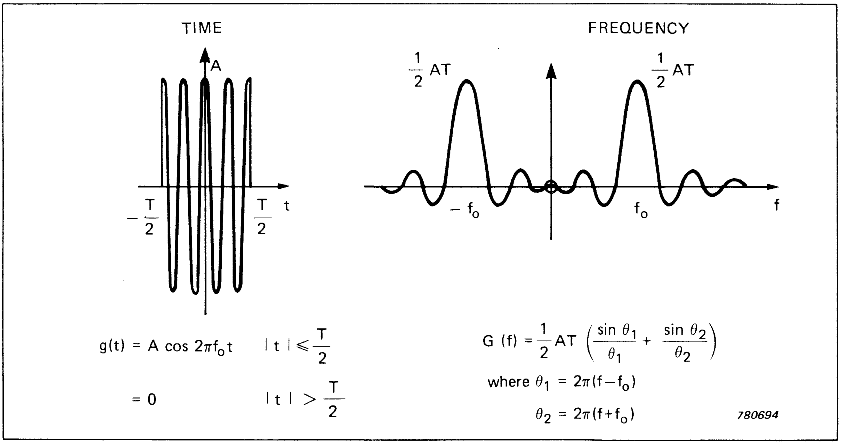

Fig.5. Frequency spectrum of a time limited sinusoidal time signal

quencies of the original sinusoid. Although it can be said that in practice we only measure the positive frequency part of this spectrum, it should be noted that the sidelobes of the "negative" sinx/x actually penetrate into the positive part and interfere with the "positive" sidelobes. This interference between the two sidelobe patterns is rather complicated and will depend on both phase and frequency of the sinusoid. We shall discuss this effect later in terms of discontinuities.

However, when we use the discrete transform we do not observe the continuous spectrum of Fig.5 but only samples of it. This is shown in Fig.6a and b. When the time limited signal contains exactly an integer number of periods the frequency samples will be taken in the centre of the mainlobe and in all the zero's between the sidelobes. Hence, we measure the true peak value of the spectrum — no picket-fence effect on the amplitude. Nevertheless, this example might be the best illustration of the picket-fence effect. The sidelobes - they are there - but behind the pickets. What we can see through the fence is just the zeros.

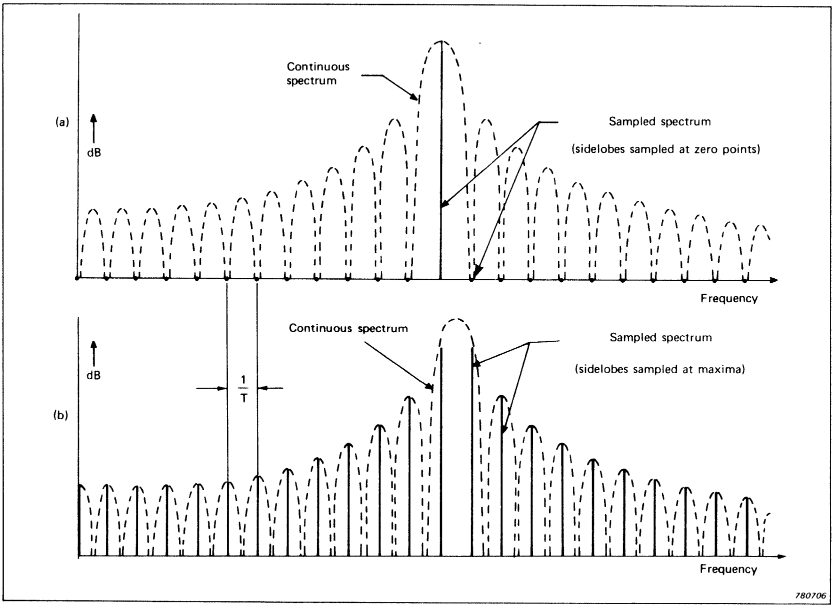

Fig.6. Frequency sampling of the continuous spectrum of a time limited sinusoid. Number of periods within the time window: a) integer, b) half-integer

This is indeed a very special situation. In the more general case where the time limited signal contains a non-integer number of periods, samples will be taken within the sidelobes and their presence will appear in the analysis. In Fig.6b this situation is shown for a half-integer number of periods within the time window. Now the sidelobes are sampled at their maxima, and the mainlobe contains two samples of equal magnitude, but both lower than the peak value. In this situation the picket-fence effect causes the maximum error in amplitude determination.

We shall now discuss this effect with real analysis examples, as shown in Fig.7. We first consider the situation of an integer number of periods, Fig.7a. Using the flat weighting we obtain the ideal spectrum having only one line at the original frequency, while the use of Hanning weighting results in three lines, indicating the broader bandwidth of the Hanning window. This is the situation of Fig.6a. These results are easily understood if we remember that the signal being analysed in an FFT analyser is not the signal in the memory but this signal repeated in all time, as a periodic signal. In this particular example, where we have an integer number of periods in the memory, the repetition of the signal will produce a perfect infinite sinusoid, from which any trace of the time limitation has disappeared. Hence, the perfect line spectrum, when we use the flat weighting. However, when the Hanning window is used, the modulation of the perfect sinusoid causes a change of the spectrum. Notice that the peak amplitudes are the same in the two cases.

Now, the examples with a half integer number of periods in the time window. Fig.7b shows a situation where the phase has been chosen so that the signal starts and stops at zero. When this signal is repeated in time we do not obtain a perfect sinusoid. Although the signal is continuous at the joints there is a change in slope. Furthermore, the signal will have a DC-component. Accordingly, the frequency spectrum obtained using the flat weighting now shows the samples within the sidelobes, as in Fig.6b. At low frequencies (DC) we find high amplitudes, while at high frequencies the amplitudes are relatively low, due to the continuity of the time signal. Also notice the picket-fence effect which has reduced the maximum amplitude by 3,9 dB. When the Hanning weighting is applied, the modulated signal will be zero and have zero slope at the joints and very low DC-level. This is also seen from the frequency spectrum, which clearly shows the improved filter characteristic of the Hanning window, and the reduced amplitude error of only 1,5 dB.

Time Signals

a)

b)

c)

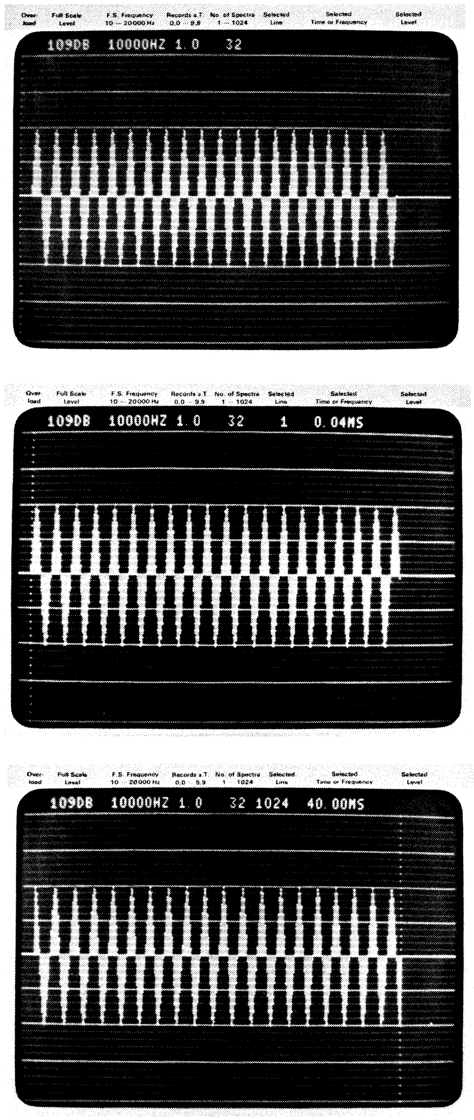

Fig.7. FFT analysis of sinusoidal time signals using different window functions. Number of periods in time memory: a)integer, b) and c) half-integer but different phases

Rectangular Weighting   Hanning Weighting

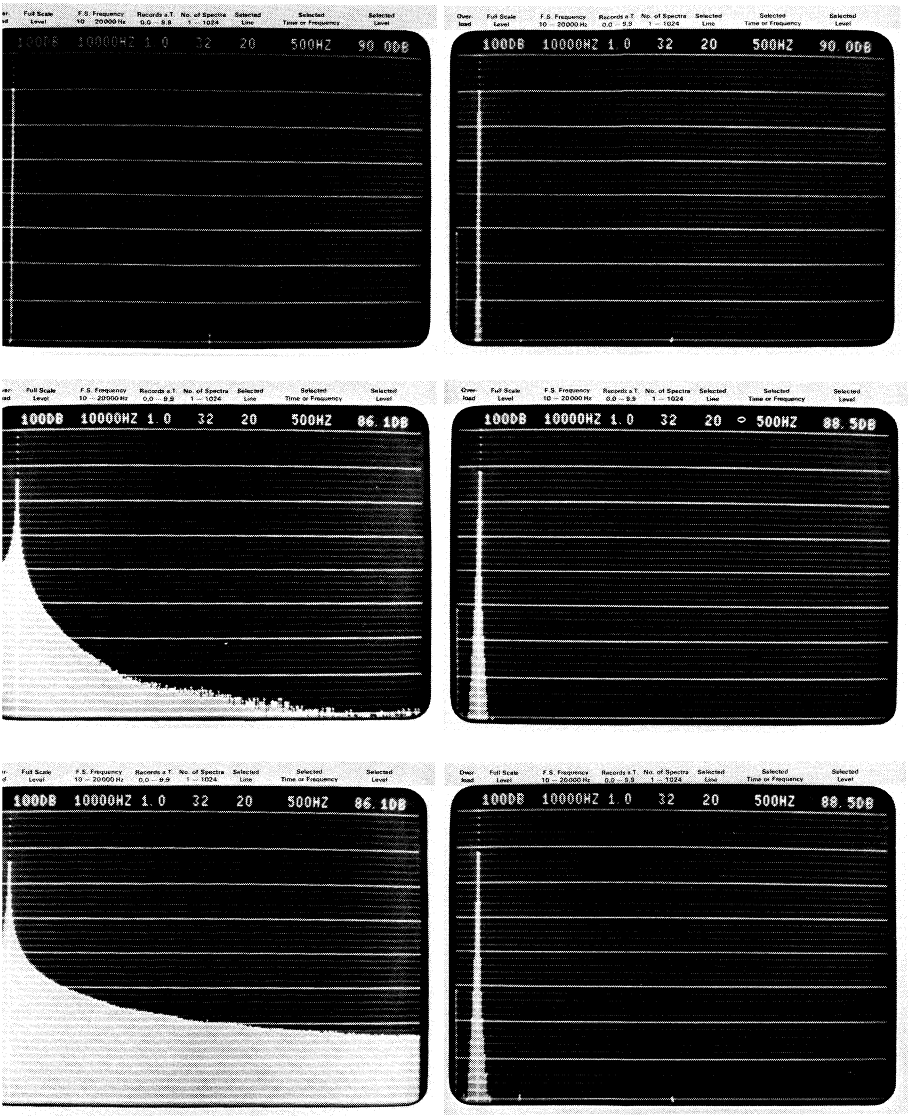

Fig. 7   Continued

In Fig.7c the same signal has been analysed, but with a different phase. The frequency is the same as in Fig.7b. When this signal is repeated in time we will find a maximum discontinuity at the joints, but now the signal will not have a DC-component. This is also shown in the analysis with the flat weighting. At low frequencies we find low amplitudes - no DC-component - while at high frequencies the discontinuity at the joints is responsible for the high amplitudes found here. Again, using the Hanning window the modulated signal will not be much different from that of Fig.7b and the analysis gives essentially the same result.

In conclusion it can be said that the purpose of the Hanning window is to remove the effect of discontinuities at the joints, and hence the Hanning weighting should be used for analysis of continuous signals, or more generally: for analysis of signals which are longer than the time window and therefore can cause discontinuities.

Of course, this situation is not special for FFT analysers. The same ar-

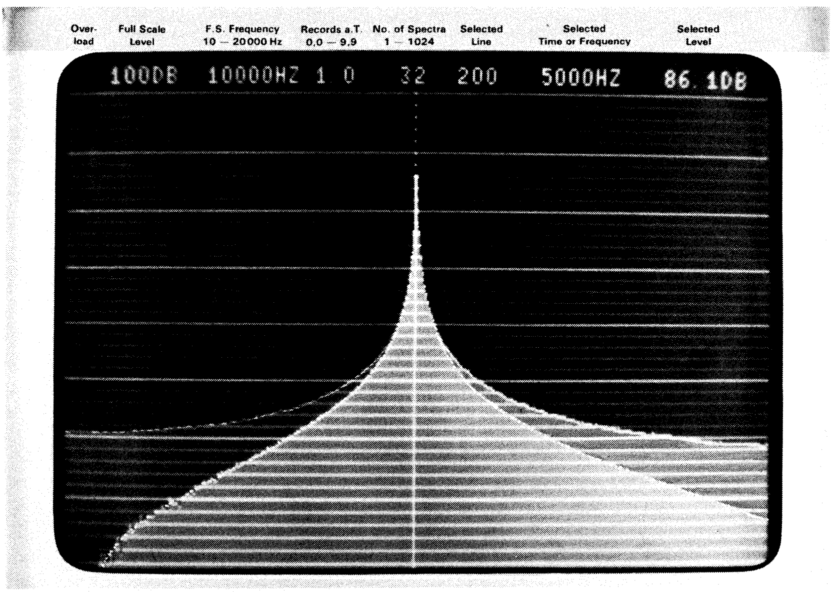

Fig.8. FFT analysis of sinusoidal signal with half-integer number of periods in memory and varying phase. The spectra show the limits for the oscillating wings, using the flat window

guments apply to Time Compression Analysers and to analog analysis of signals recorded on a magnetic tape loop. Only the way the effect shows up is more pronounced with FFT analysers. If, e.g. the signal of Fig.7b and c were analysed with a phase changing with time, i.e. without any synchronisation between the signal and the recording of it, the use of the flat window would result in a spectrum where the "wings" would move up and down with time, maybe giving associations of a "Hawaian hula dance". In Fig.8 is shown the limits of these wing-oscillations, by superposition of the two extreme cases.

But the Flat weighting function - when should it be used? The answer is: for transient signals which can be completely contained within the time window. In this case the signal itself starts and stops at zero, and there will be no discontinuities when it is repeated in time. So there is no need for a Hanning window. But more important, if the Hanning window was used it would give different weighting to different parts of the transient and therefore result in a wrong analysis. This is illustrated in Fig.9, where a transient signal - being just one period of a sinusoid - is analysed at different positions within the time window. Using the Flat window we obtain the correct analysis, which is independent of the position — as it should be (the difference between the two spectra is caused by different averaging —over 1 and 64 individual spectra respectively). This is not the case when the Hanning window is used. In Fig.9a the transient is situated in the beginning of the time window where the Hanning weighting gives a high attenuation, which also changes over the length of the signal. Accordingly, the spectrum has too low amplitudes and is somewhat distorted. Linear averaging over 64 spectra was used in this case to decrease the influence of noise. Fig.9b shows the situation where the signal is situated at the centre of the Hanning window. Since the weighting function is rather flat here, only little distortion might be introduced, but the amplitudes are now ≈6 dB too high. As mentioned previously, this is due to the amplitude of the Hanning window being 2 at its centre.

It should be mentioned that transient analysis is normally done in terms of energy or energy density. However, the spectra produced by FFT analysers are power spectra, since they are analyses of periodic signals. In order to convert from power to energy units it is therefore necessary to multiply with the length of the time window, being the repetition time of the periodic signal.

Since FFT analysers perform the analysis on blocks of time data this makes them extremely well suited for the analysis of transients. The

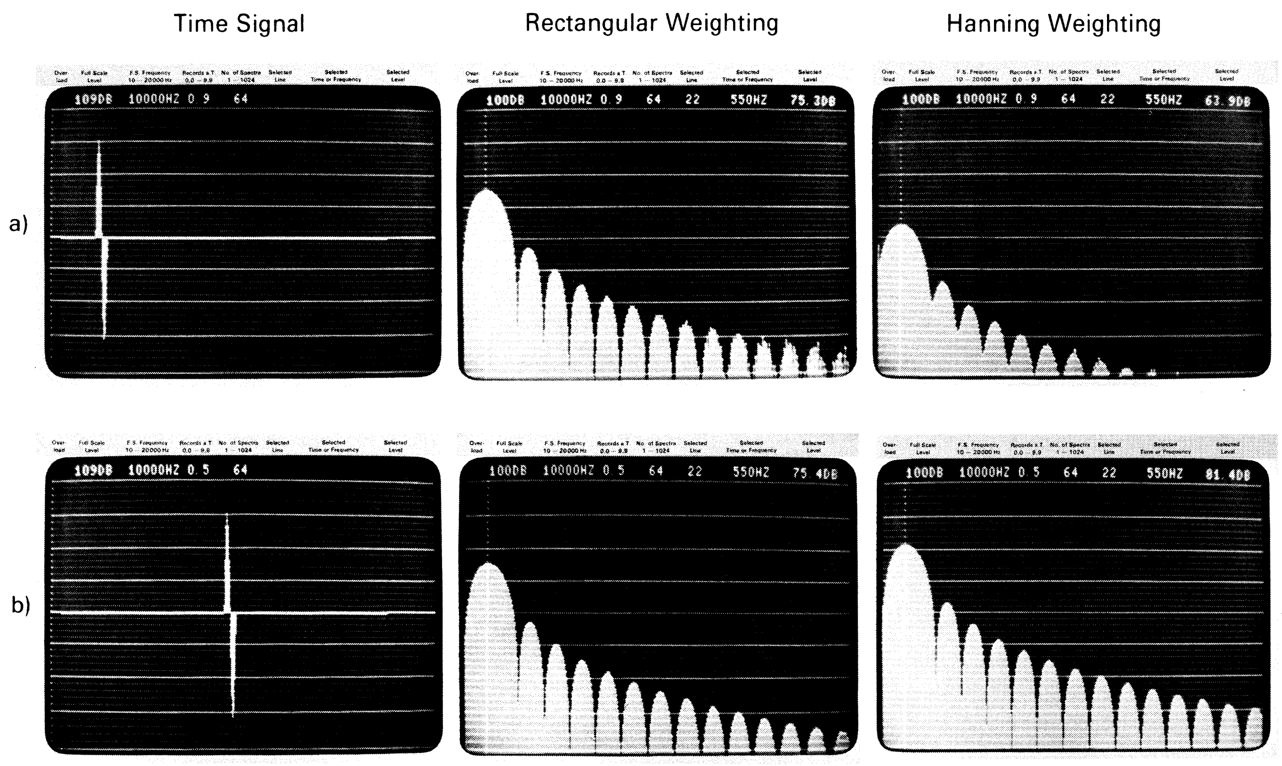

Fig.9. FFT analysis of a transient signal positioned a) at the beginning of the time window, b) at the centre of the window

only problem involved is the capture of the signal within the time window. However, since the FFT technique is purely digital it is rather simple to provide the analyser with transient recording facilities, like a trigger function - as found on oscilloscopes - but also with an adjustable delay between the trigger point and the recording. It is interesting that even negative delays, i.e. pre-trigger recording, can also be achieved. In this way FFT analysers are extremely easy to operate with respect to transient analysis.

To make a very short conclusion of this rather long discussion:

Hanning weighting is for long continuous signals.

Flat weighting is for short transients.

The FFT algorithm

It was mentioned in the beginning of this paper that FFT analysers worked on the basis of the Discrete Fourier Transform, but calculated this in a very efficient way by use of the Fast Fourier Transform, the FFT-algorithm. Hence, FFT is not a Fourier transform in itself, but merely a calculation scheme. From an application point of view it is of little interest to know how the spectra are calculated, with or without FFT, since the results are the same. Actually the use of the FFT algorithm only influences one specification of an analyser - its real-time performance - since this will depend directly on the calculation time.

We shall therefore only discuss the FFT-algorithm very briefly.

The Discrete Fourier Transform takes the form shown in Fig. 1d:

$$G(k) = \frac{1}{N} \sum_{n=0}^{N-1} g(n) \cdot e^{-j \frac{2\pi kn}{N}}$$

This equation can be written as a matrix equation:

$$\{G_k\} = \frac{1}{N} \{A_{kn}\} \{g_n\}$$

where $\{G_k\}$ and $\{g_n\}$ are column vectors containing the N frequency samples and the N time samples, respectively. $\{A_{kn}\}$ is a square matrix of order N containing the complex unit vectors, $e^{-j2\pi kn/N}$

For the particular case of N = 8 the matrix equation may be visualised as follows:

$$\begin{bmatrix} G_0 \\ G_1 \\ G_2 \\ G_3 \\ G_4 \\ G_5 \\ G_6 \\ G_7 \end{bmatrix} = \frac{1}{8} \begin{bmatrix} \uparrow & \uparrow & \uparrow & \uparrow & \uparrow & \uparrow & \uparrow & \uparrow \\ \uparrow & \nearrow & \rightarrow & \searrow & \downarrow & \swarrow & \leftarrow & \nwarrow \\ \uparrow & \rightarrow & \downarrow & \leftarrow & \uparrow & \rightarrow & \downarrow & \leftarrow \\ \uparrow & \searrow & \leftarrow & \nearrow & \downarrow & \nwarrow & \rightarrow & \swarrow \\ \uparrow & \downarrow & \uparrow & \downarrow & \uparrow & \downarrow & \uparrow & \downarrow \\ \uparrow & \swarrow & \rightarrow & \nwarrow & \downarrow & \nearrow & \leftarrow & \searrow \\ \uparrow & \leftarrow & \downarrow & \rightarrow & \uparrow & \leftarrow & \downarrow & \rightarrow \\ \uparrow & \nwarrow & \leftarrow & \swarrow & \downarrow & \searrow & \rightarrow & \nearrow \end{bmatrix} \begin{bmatrix} g_0 \\ g_1 \\ g_2 \\ g_3 \\ g_4 \\ g_5 \\ g_6 \\ g_7 \end{bmatrix}$$

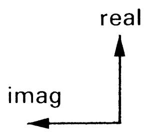

Each of the arrows in the matrix represents a complex unit vector, relative to the axes shown. It is seen that a direct calculation of this matrix equation would require N^2 complex multiplications, and when considering calculation time, multiplications are the most time consuming operations. Using the FFT algorithm the number of multiplications can be reduced to $N\log_2 N$, assuming N is a power of two. In the typical case of N = 1024, the reduction in calculation time is more than 100.

The savings of the FFT algorithm result from a factorisation of the matrix into a number ($\log_2 N$) of matrices. Actually it is not the matrix $\{A\}$ which is factorised but a different version of it, matrix $\{B\}$ shown in Fig.10. $\{B\}$ is derived directly from $\{A\}$ by interchanging rows with so-called "bit-reversed" numbers, also indicated in Fig.10. It can be shown that the matrix $\{B\}$ can be factorised into the three ($\log_2 8 = 3$) matrices $\{X\}$, $\{Y\}$ and $\{Z\}$ also shown in Fig.10. It is a common property of these three matrices that each row only contains two factors different from zero, one of which is always unity. Hence multiplication with one of these matrices requires only N multiplications, one for each row. Multiplying all the matrices thus requires $N \log_2 N$ multiplications. Due to the interchange of rows in $\{B\}$ the final result will contain the frequency components in the interchanged positions. However, putting the components back in their correct positions is a very fast procedure compared to the matrix multiplications.

But it is possible to speed up the calculations even further. Notice that in matrix $\{Z\}$ rows separated by 4 rows are identical except for a change of sign of the non-unity factor. Hence, such two rows can be calculated using only one multiplication, by changing an addition to a

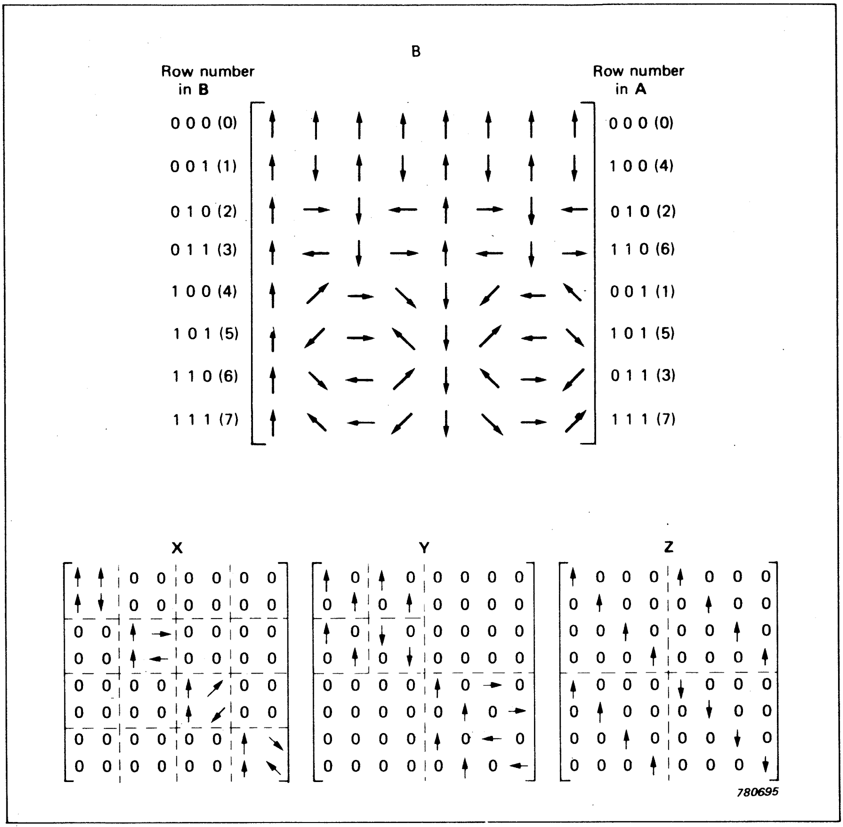

Fig. 10. Factorisation of matrix B used in the FFT algorithm.
{B} = {X} · {Y} · {Z}

subtraction. In {X} and {Y} the separations between these nearly identical rows are 1 and 2 respectively. In this way an additional factor of two can be saved in calculation time. Also, since the time signal to be transformed is normally a real function while {g} vector is complex, a more efficient transform can be made by packing the N-point real time signal into a N/2-point complex signal. Hence an N-point transform can be performed using only a N/2-point calculation. Of course, such a mixing up of the time signal will require some additional calculations after the transform to obtain the frequency spectrum of original signal. However, time is still saved.

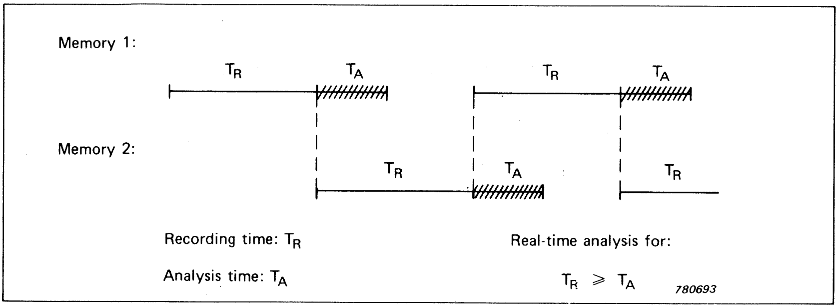

Fig. 11. Time chart showing time used for recording and analysis of signals in FFT analysers

The calculation time has direct influence on the real time performance of an FFT analyser. This is shown in Fig. 11. In order to obtain real-time analysis, all of the time signal should be analysed, without any loss of data. Since the analyser first records a block of time data and then performs the analysis of this block, it is necessary to work with two time memories. While calculations are performed on the data in one memory, further time data can be recorded in the other. It is obvious that in order not to lose any time data the time used for recording must be larger than the time used for the analysis. The recording time depends on the frequency range, i.e. on the sampling frequency, since the number of time samples in the memory is constant. Hence, there will exist an upper frequency limit above which the analysis cannot be performed in real time. This real-time frequency will in practice be determined by the speed of the electronic circuits used in the design, and will therefore directly influence the cost of the analyser. It is therefore of importance to determine whether a specific measurement requires real-time performance or not. A thorough discussion of this subject will be found in reference [5].

Conclusions

The advantages of FFT analysers are many: very good linearity of both frequency and amplitude, large dynamic range, extreme stability, easy analysis of transients, linear averaging, reference memories, spectrum comparisons, etc. etc. Furthermore, they are easy to connect to other types of digital equipment, such as calculators, allowing for a further treatment of the data - even automatic analysis procedures can be controlled easily.

So the FFT analyser is a very flexible instrument, but still easy to operate. It will deliver analysis results with a speed of several spectra per second, results which are reliable and can be reproduced. However, there are a few pitfalls. But these can easily be avoided when one knows the assumptions on which the Discrete Fourier Transform works and the approximations involved.

References

[1] A. PAPOULIS: The Fourier Integral and its Applications. McGraw Hill Inc., New York, 1962

[2] E. O. BRIGHAM: The Fast Fourier Transform. Prentice Hall Inc., New Jersey, 1974

[3] R. B. RANDALL: The Application of B & K Equipment to Frequency Analysis. Brüel & Kjær, 1977

[4] R. UPTON: The Objective Comparison of Analog and Digital Methods of Real-Time Frequency Analysis. B & K Technical Review, No.1, 1977

[5] R. B. RANDALL and R. UPTON: Digital Filters and FFT Technique in Real Time Analysis. B & K Technical Review, No.1, 1978

News from the Factory

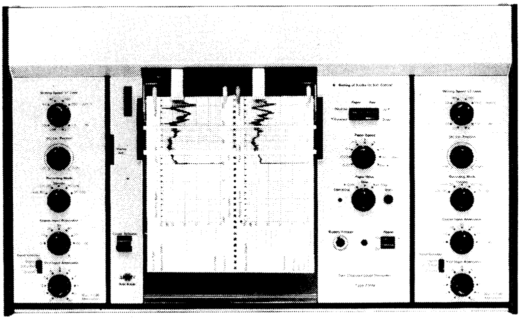

Two-Channel Level Recorder Type 2309
Graphic recording in the field as well as in the laboratory is an invaluable means for providing a hard copy for immediate documentation of the measured results. The Two Channel Recorder Type 2309 is a portable, battery operated, compact instrument for plotting two varying parameters simultaneously with respect to a third, for example, time, frequency, or r.p.m.

Each channel of the 2309 has four switch-selected recording modes: "AC Log", "DC Log", "—DC Lin", "+DC Lin" and a "Stand by" mode where the signal to the writing system is switched off. In the "AC Log" mode it records the RMS value of any waveform in the frequency range 1,6 Hz to 20 kHz within an accuracy of ±0,5 dB for signals with crest factors up to 3. In the "DC Log", "—DC Lin" and "+DC Lin" modes the signal from the range potentiometer is fed through a 500 Hz electronic chopper before it enters the RMS rectifier circuit. In the DC Lin modes the writing system is able to follow a sinusoidal signal with a maximum frequency of 1,6 Hz with maximum deflection and at the fastest writing speed.

The dynamic range of the Level Recorder is determined by the interchangeable logarithmic potentiometer of 25 dB or 50 dB supplied for both the channels. In the Lin recording mode an antilogarithmic amplifier is inserted in the signal path just before the range potentiometer.

Each channel has four writing speeds, 16, 40, 100 and 250 mm/s corresponding to lower frequency limits of 1,6, 4, 10 and 25 Hz. Recordings are made on 50 mm wide preprinted frequency-calibrated paper as a function of frequency or on lined paper as function of time or other parameter. Eight switch-selected paper speeds from 0,01 mm/s to 30 mm/s are available and the start and stop of the paper drive can be remotely controlled. Additionally, the recorder has a wide range of facilities for filter, generator, and analyzer synchronization and external control and synchronization of the paper movement.

The power for driving the recorder is delivered from dry cells or rechargeable NiCd-cells mounted in a plug-in battery box, from mains via the optional plug-in Power Supply Type 2808 or from an ordinary automobile battery. A built-in miniature meter on the front panel is also provided for monitoring the supply voltage.

Among the numerous applications of the instrument are recording of phase and distortion responses of stereo equipment; phase and amplitude responses of filters, amplifiers etc.; sound insulation of building elements with respect to frequency; excitation force and vibration response as a function of frequency for mechanical structures; and vibration levels and other parameters as a function of machinery rotation speed.

Bump Recorder Type 2503

Goods items during transport are often exposed to higher shock levels than they are designed to withstand. By monitoring the severity of shocks with respect to time using a Bump Recorder Type 2503, it is possible to estimate whether damage is likely to have taken place during transport and to indicate where the responsibility lies from the time of occurrence.

The shocks are picked up by a triaxial accelerometer which can be mounted either inside the recorder case or on a critical part of the transported item. The accelerations are monitored in three mutually perpendicular directions so that shocks occurring in any random direction are resolved into three vectors. These are combined in the 2503 to represent the magnitude of the applied shock.

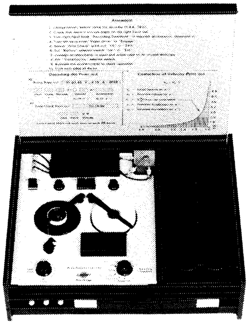

A "Recording Threshold" control is provided to limit the data print out to shock events which have a severity approaching an estimated serious level. The print-out is only activated when this level, which can be preset between 10 and 100 m/s^2 (1 — 100 g), is exceeded. Bump data is recorded on a 6 mm wide paper strip, first the day, hour and minute of occurrence then the maximum velocity and acceleration levels. Built-in rechargèable batteries power the instrument for approximately 18 days, for longer journeys an external battery pack can be connected.

In the design and development of packaging, the Bump Recorder will provide valuable data by measuring the shock levels experienced by a product inside a container under various drop conditions.

By using a low g accelerometer Type 8306 with the 2503, shock levels as low as 0,01 m/s^2 (0,001 g) or 0,3 mm/s can be recorded. The ability to record such low levels enables the Bump Recorder to be used for monitoring ground vibrations due to heavy traffic, explosive detonations, pile driving, seismic activity etc.