Research and Development Report

CONTROLLED IMAGE DESIGN:
The measurement of Time-Frequency responses

R. Walker, B.Sc.(Eng), C.Eng., F.I.O.A., M.I.E.E.

Research and Development Department
Technical Resources Division
THE BRITISH BROADCASTING CORPORATION

CONTROLLED IMAGE DESIGN:
The measurement of time-frequency responses

R. Walker, B.Sc.(Eng), C.Eng., F.I.O.A., M.I.E.E.

Summary

This Report describes the measurement of acoustic events in a 3-dimensional measurement domain. An outline of the general theoretical background is followed by a description of the special requirements for the measurement of short-term acoustic responses in rooms. The results of measurements on an experimental synthesis of a single room reflection are also presented. It is shown that the measurements and presentation of the result in terms of amplitude/time, amplitude/frequency and 3-dimensional time/frequency/amplitude responses accurately portray the true situation, within the theoretical limitations of the Fourier Transform. It is shown that the achievable time and frequency resolutions are probably just adequate for the measurement of those effects thought to be important for the perception of the stereophonic illusion.

Issued under the Authority of

Research & Development Department,
Technical Resources Division
BRITISH BROADCASTING CORPORATION

General Manager
Research & Development Department

CONTROLLED IMAGE DESIGN:
The measurement of time-frequency responses

R. Walker, B.Sc.(Eng), C.Eng., F.I.O.A., M.I.E.E.

CONTROLLED IMAGE DESIGN: THE MEASUREMENT OF TIME-FREQUENCY RESPONSES

R. Walker, B.Sc.(Eng), C.Eng., F.I.O.A., M.I.E.E.

1. INTRODUCTION

Since the work on the effects of early reflections in control rooms began[1], the need has arisen to measure a number of parameters of the acoustic signal in short intervals of time. The conventional way of describing the characteristics of acoustic or audio systems is as functions of frequency.

It is well known that the system impulse response theoretically contains all of the information necessary to specify a system fully, at least for a time-invariant one. However, the impulse response is purely a time-domain function (although it can be converted to other forms of representation by suitable mathematical operations). It may or may not effectively be frequency limited, depending on the equipment used to measure it, but there is no direct information in the impulse response to indicate the frequency content.

Many years ago, Shorter[2] used time-domain gating followed by conventional frequency response measurement to find the short-term response of loudspeakers. More recently, but still some years ago, the present author used a similar method to measure the time response of rooms[3]. In both cases, the methods were relatively clumsy, inflexible and extremely laborious to apply; to the extent that no great use was made of them in either case.

Modern methods are simple to apply, using commercially available equipment. Some proponents of such measurement methods have succeeded to the point where their methods are accepted as *de facto* standards for such measurements[4].

The recent revival of interest in these types of measurements arose mainly from the Controlled Image Design (CID) work. This work sought to improve the quality and consistency of the listening conditions in studio control rooms by reduction of the spurious early reflection energy. The principle relies on providing angled wall and ceiling surfaces to direct the unwanted reflected sound energy away from the main listening position[5,6]. By this means, the distortions of frequency response and stereophonic image quality caused by the early energy could be effectively eliminated. The objectives were not only to provide a more accurate monitoring quality but also to produce more consistency between different rooms.

It was important that the means of measurement used in the development of the prototype CID room and in the checking and correction of practical implementations of the principle was properly validated. This Report presents the theoretical and practical backgrounds to the measurement process, together with the inherent theoretical limitation. It also presents the results of measurements carried out using artificially-generated acoustic reflections, both wide-band and filtered.

2. THE FOURIER TRANSFORM

All of the useful interpretations of time-domain events in the frequency-domain (and indeed, by implication, most aspects of analogue circuit theory) are based on the Fourier Transform. This provides a means of translating between the frequency and time domains.

The theoretical basis of the Fourier Transform can be found in any text book on circuit theory or signal processing. In principle, the whole infinite time record is convolved with an infinity of in-phase and quadrature sinusiodal components to obtain an infinite series of (complex) components. It can be shown that the original time signal and the summation of the sinusiodal components are identically equivalent representations*.

In the limit, for perfect frequency resolution, the signal must exist, and be available to the analysis, for all time (–inf to +inf). Conversely, an infinitesimal time event carries no frequency information at all.

In the real world, such theoretical signals cannot exist. Ideally, even a measurement of a sinusoidal waveform from a simple oscillator requires infinite time. In practice, it is clearly sufficient to limit the time-domain record to some reasonable length, such that the effects

* One example which often causes confusion is the concept of ‘transients’ in audio signals. Most so-called ‘audiophiles’ recognise that an electronic circuit has to respond faithfully to the rapidly-changing time-waveforms of real audio signals. Most also understand the concept of an audio bandwidth (even if it is sometimes only very roughly commensurate with that perceptible to human beings). Few recognise the inescapable link between them. Some years ago it became fashionable to consider a certain type of distortion in audio amplifiers which, in any reasonable amplifier, can only be caused by applying transient signals outside the working range of the amplifier. It can be entirely prevented by the use of a suitable low-pass filter. Another example widely misunderstood is the purpose and behaviour of loudspeaker crossover filters. In this case, the complicated time-function of the audio signal is split into different frequency bands (that is, processed by a kind of real-time, analogue Fourier Transformer), to be recombined acoustically; the objective, in principle, being to reproduce the original time-domain waveform. Even some audio professionals think that the low-frequency path (that is, the ‘woofer’) has to respond rapidly to transient signals and are puzzled as to how that can be achieved, given the great mass and slow response time of the low-frequency drive unit.

of the truncation are acceptable, in the context of the measurement.

In effect, the product of time and frequency resolutions is a constant, approximately equal to unity. Intuitively, it is obvious that information must be available for a reasonable fraction of a full cycle before anything reliable can be inferred about the frequency or phase of a sine wave. Thus, to make statements about time intervals of the order of a few milliseconds implies frequency resolutions of around 1 kHz.

The usefulness of the Fourier Transform for the description of short time events lies in the fact that Fourier transformation of the impulse response gives the frequency-domain response called the Transfer Function — equivalent to the response of the system, in amplitude and phase, to excitation at each frequency separately by a sinusoidal waveform which has existed for all time. This is what is commonly understood to be the steady-state frequency response.

3. DISCRETE FOURIER TRANSFORMS

The theory outlined above applies to a time-domain signal of infinite duration and which is continuous, that is it is defined at every instant. The resulting transform has an infinity of components with infinitesimal frequency-domain spacing. None of these aspects is very practical. The time domain can be limited to just one period if the signal is assumed to be repetitive. Then the frequency-domain components exist only at multiples of the cyclical frequency (i.e. at multiples of 1/period). They still extend up to infinite frequency, but the higher ones (and the negative ones) can always be ignored in practice; that is, above the required system frequency limit. In a real, band-limited system they would be negligible anyway. Theoretically, complications arise if the amplitude distribution is not continuous but is quantised, as in digital systems. For practical purposes, this introduces additional noise and can be ignored if the resulting dynamic range of the measurement system is adequate for the purpose. Most modern measurement systems are equivalent to at least 12-bit resolution, corresponding to reasonably accurate representation of components down to about −60 dB.

If the time axis is also quantised, as it will be in any measurement system based on sampling and digital processing (as most are), then a further restriction applies to the data and its Fourier Transform, which then becomes a Discrete Fourier Transform (DFT). According to the normal (Nyquist) sampling theory, the highest frequency component which can be represented correctly is half of the sampling frequency. If the continuous signal is not pre-filtered, higher frequency components will be aliased into the baseband and will not be distinguishable from the real baseband signals. This is not usually a problem because anti-alias filters are always included in any realistic measurement system. Minor difficulties may be encountered with the residual alias components in some cases, especially if the time signal contains constituents which would otherwise transform to high levels at frequencies outside the range of immediate interest. This is especially true in the region close to the Nyquist frequency.

4. TRANSFORM WINDOWS

A continuous, non-repetitive signal (for example, the whole or part of an impulse response) can be treated as repetitive, and therefore amenable to transformation to the frequency domain, if part of it is selected and assumed to be repetitive. In the frequency domain, the associated loss is that of frequency resolution. The transform will only contain information about frequencies harmonically related to 1/period. The manner of selection also has important consequences. The weighting function used to select the data in the time domain is called a 'window'.

The theory of data truncation and windows can be found in any text book on signal processing. (Ref. 7 is a particularly comprehensive review and analysis of 44 different windows.) In brief, the frequency-domain response obtained from the transform includes convolution with the transform of the window. In general, this will have a main response of finite width and a sidelobe response extended to frequencies quite remote from the centre. As an example, Fig. 1 shows the transform of a rectangular window. It has a first sidelobe amplitude of −13.3 dB and subsequent sidelobes which reduce at a rate of 20 dB per decade. This response means that each component in the Transfer Function may contain quite large contributions from nearby frequencies and some contributions from quite remote frequencies, depending on the shape of the window in the frequency domain. This is known as 'leakage'.

A large number of windows have been developed over a period of many years to optimise the response for different purposes. Table 1 *(see page 4)* shows the main characteristics of some of them[8].

If a signal really does exist in time for only a short period and is isolated from other signals then it is comparatively simple to decide which part to select. The effective window can be rectangular, with a finite value during the signal, zero outside and infinitely sharp transitions between the two regions. Apart from the inevitable limitation of frequency resolution, nothing else is lost (or gained).

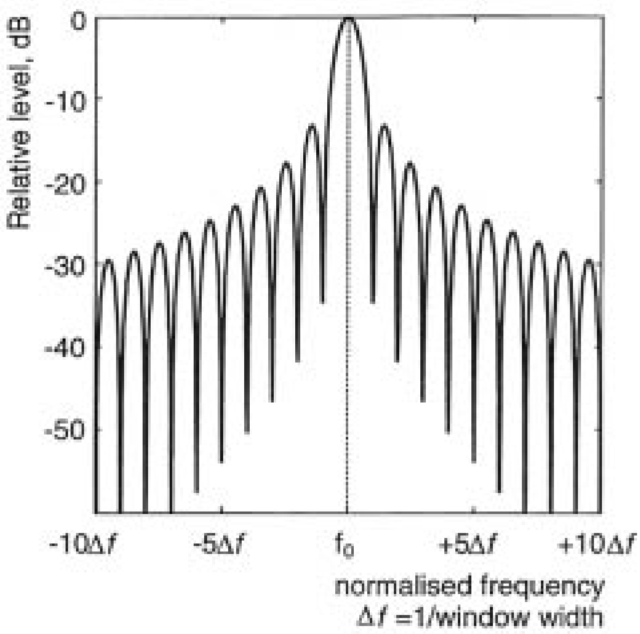

Fig. 1 - Fourier transform of rectangular time-domain window.

When the signal also exists outside the time zone of interest the application of a rectangular window would create alias components by the resulting abrupt transition at the edge of the window. These alias components would contaminate the wanted signal because of the leakage corresponding to the sidelobe response of the window. Thus, some form of tapered transition is necessary to reduce such effects to obtain a better compromise between frequency discrimination and spurious response.

Many forms of gradual transition windows have been developed for different purposes. As shown in Table 1, *(overleaf)* these window functions have different properties; some produce high resolution at the expense of leakage and others give high rejection of sidelobe response, at the expense of poorer frequency resolution. A gaussian window is sometimes used. This has the same gaussian shape in both time and frequency domains and, therefore, has no sidelobe response at all (in either domain), at the cost of a rather poor frequency resolution.

5. ACOUSTICAL TIME-FREQUENCY RESPONSES IN ROOMS

5.1 Psycho-acoustic properties of hearing

The human hearing system is a complicated signal processor, especially in the context of the effects involved in the creation of a stereophonic audio illusion. The following is a very rough summary of a great deal of psycho-acoustic research, stated without references because of the large and sometimes contradictory body of evidence.

In the interval after the arrival of the direct sound (up to about 5 ms), any delayed sound is integrated with the direct sound to form some impression of the sound ‘quality’. The sense of direction is governed by the arrival of the direct sound.

From about 5 ms to about 20 ms, information is extracted about the direction of a sound source. After about 20 ms, the extracted information is largely about the surrounding space rather than the source itself, until about 50 – 80 ms when the sound becomes distinctly reverberant.

5.2 Acoustical time-frequency responses in typical control rooms

For reflections in a (small) room, the interval up to about 15 – 20 ms is the most important. Although in the period up to 5 ms reflections are not perceived directly, they can have an important influence on the sound quality. A relative time delay of 1 ms, for example, could produce a strong cancellation by interference with the direct sound at multiples of 500 Hz. In most cases of very short delays, a surface would not be large enough to cause a strong reflection at 500 Hz, but it might be at 1500 Hz and higher harmonics. Such cases are frequently encountered in control rooms, where the flat surface of the mixing desk usually forms an efficient reflector. The (potential) reflection from the top surface of the mixing desk is likely to occur at about 0.8 to 1.2 ms after the direct sound from the loudspeakers.

Reflections later than 5 ms can disturb the subjective stereophonic imaging process, causing mislocation of images. Individual early reflections from room surfaces are likely to occur at about 3 ms (from the ceiling), 7 – 8 ms (from the side walls) and 15 – 20 ms (from the rear wall). Thus, for the measurement of the acoustical effects of early reflections in control rooms, it is desirable to be able to resolve time differences of the order of 1 ms.

Fortunately, the stereophonic illusion principally involves the higher frequencies. Although some sense of spaciousness is conveyed by frequencies between about 300 and 1000 Hz, the main image-forming frequencies are those from about 1000 Hz upwards. It is desirable to obtain as high a frequency resolution as possible, implying longer time records, in order to obtain some idea of the frequency character istics of any reflections.

It is also important to identify isolated reflections amongst others, so that the sideband rejection is also required to be high.

*Table 1: Frequency domain parameters of some window types**

Window	Noise Bandwidth	3 dB Bandwidth	Ripple dB	Highest sidelobe dB	Sidelobe fall off dB/dec	60 dB Bandwidth
Rectangular	Δf	$0.89\Delta f$	3.92	–13.3	–20	$665\Delta f$
Hanning	$1.5\Delta f$	$1.44\Delta f$	1.42	–31.5	–60	$13.3\Delta f$
Kaiser-Bessel	$1.8\Delta f$	$1.71\Delta f$	1.02	-66.6	–20	$6.1\Delta f$
Flat-top	$3.77\Delta f$	$3.72\Delta f$	0.01	-93.6	0	$9.1\Delta f$
Blackman-Harris**	$2.00\Delta f$	$1.90\Delta f$	0.83	–92.0	–6	$\approx 6\Delta f$

where Δf is the frequency domain line spacing (= 1/period)

Rectangular:

$w(t) = 1$ for $0 <= t < T$,
$w(t) = 0$ elsewhere.

Hanning:

$w(t) = 1 - \cos 2\pi t/T$ for $0 <= t < T$,
$w(t) = 0$ elsewhere.

Kaiser-Bessel:

$$w(t) = 1 - 1.24 \cos 2\pi t/T + 0.244 \cos 4\pi t/T - 0.00305 \cos 6\pi t/T \quad \text{for } 0 <= t < T,$$

$w(t) = 0$ elsewhere.

Flat-top:

$$w(t) = 1 - 1.93 \cos 2\pi t/T + 1.29 \cos 4\pi t/T - 0.388 \cos 6\pi t/T + 0.0322 \cos 8\pi t/T \quad \text{for } 0 <= t < T,$$

$w(t) = 0$ elsewhere.

Blackman-Harris[8]

$$w(t) = 0.35875 - 0.48829 \cos 2\pi t/T + 0.14128 \cos 4\pi t/T - 0.01168 \cos 6\pi t/T \quad \text{for } 0 <= t < T,$$

$w(t) = 0$ elsewhere.

* From Ref. 7 (except for Blackman-Harris).

** The ultimate behaviour of the Blackman-Harris window is particularly dependent on the accuracy of representation of the coefficients. The figures given are for the theoretical window, giving a theoretical suppression of the first sidelobe of 92 dB.

These factors lead to measurement processes based on time intervals of about 1 – 2 ms, using some form of tapered window, resulting in frequency resolutions of the order of 0.5 to 1 kHz. It is, as a result, conceptually possible to identify and measure reflections, with time and frequency resolutions high enough to be useful for the investigation of stereophonic systems in relatively small rooms.

5.3 Measurement and presentation of time-frequency responses

There are two particularly useful measures of time-domain responses — the so-called ‘Energy-Time’ response (ET) and the 3-dimensional Energy-Time-Frequency response (ETF). The first of these, the Energy-Time curve, is in fact the magnitude of the complex system impulse response. It is taken to represent ‘instantaneous’ energy. It may be followed by some form of filtering window to highlight the main features. Although its precise theoretical nature is the subject of some discussion[9], it does present a view of the time-domain response in which representations of discrete reflections are easily observed.

The second useful measure of response is the 3-dimensional, ‘waterfall’ plot. For this, the time-domain impulse response is translated to the frequency domain by the discrete Fourier transform (DFT), usually with a half-Hanning window — half, in order to reject all preceding data, at the cost of some aliasing. The start of the transform block is progressively shifted to later times, to produce a series of frequency responses at different times. It is also limited in time resolution by the length of the DFT window. Despite these limitations, a useful display indicating approximate times and frequency responses of reflections can be obtained.

One convenient instrumental system which greatly simplifies the measurement and post-processing of impulse response functions is MLSSA (Maximum Length Sequence System Analyser)[10]. This instrument comprises source signal generator and response recorder which works not on the impulse response directly but by using a comparatively long-duration, noise-like test signal. The actual impulse response is obtained by cross-correlation. This results in a significant degree of noise rejection and a more reliable measurement. The system also includes a wide range of post-processing functions. All of the measurements presented in this report were obtained using MLSSA.

Refs. 5, 6 and 11 give the results of measurements carried out, using these methods, in a number of implementations of the Controlled Image Design principles. In these real cases, significant irregularities in the frequency responses of most reflections gave rise to apparent discrepancies between the ET and the EFT responses. Overall, the reflection amplitudes achieved were low, mostly less than –20 dB relative to the direct sound. There were, however, some anomalous, isolated narrowband effects with amplitudes somewhat higher than that.

6. EXPERIMENTAL RESULTS

6.1 Overview

In the initial stages of the work on early reflections in rooms[1,5,6,11], it was noticed that the selection of filtering, transform lengths, etc. had a pronounced influence on the appearance of the results. For example, in some cases where the ET response apparently showed reflections of the order of −20 dB or lower, the EFT response showed narrowband reflections much higher than −20 dB at some frequencies. Figs. 2 and 3 *(overleaf)* show an example from Ref. 11. In that case, the reflection at about 7.5 ms appears in the ET response to be at −22 dB, whereas the EFT response shows it to be at −10 dB at 3800 Hz (the display level is such that the direct sound is at approximately +3 dB). The ET response was calculated with the Blackman-Harris window, which rejects just the extreme high and low frequencies, leaving most of the middle frequency range evenly weighted. The EFT response was measured with a 4.26 ms transform length, with a half-Hanning window, giving an effective frequency resolution of about 300 – 400 Hz.

It is clear that the wideband ET response, Fig. 2, is producing an overall average value for the amplitude and that, as the frequency resolution is increased (for the EFT response, Fig. 3), the amplitudes of both direct and reflected sound may be altered. In the case of the wideband direct sound, the measured level decreases as the bandwidth is reduced (for a fixed total energy), whereas the amplitude of a narrowband effect is relatively unaltered (if the filter is centred on the maximum response of the effect). This is also observed if the ET response is restricted to narrower bands.

6.2 Experimental synthesis of reflections

An experimental simulation of a single reflection was set up using an electronic delay line. The delay was set to 5 ms and a mixture consisting of the input signal and the output at −6 dB was taken as the test signal. A bandpass filter of nominally 4 kHz to 6 kHz was optionally inserted into the delay path. This filter actually produced a passband gain of +2 dB, making the filtered delayed signal −4 dB relative to the input signal. (Because of the simple mixer, there were other signals corresponding to multiple passes through the delay

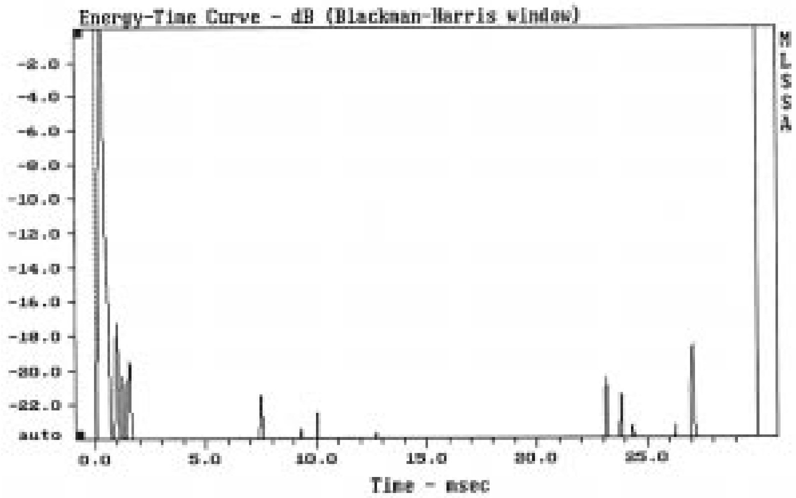

Fig. 2 - ET response, showing low reflection amplitude at 7.5 ms.

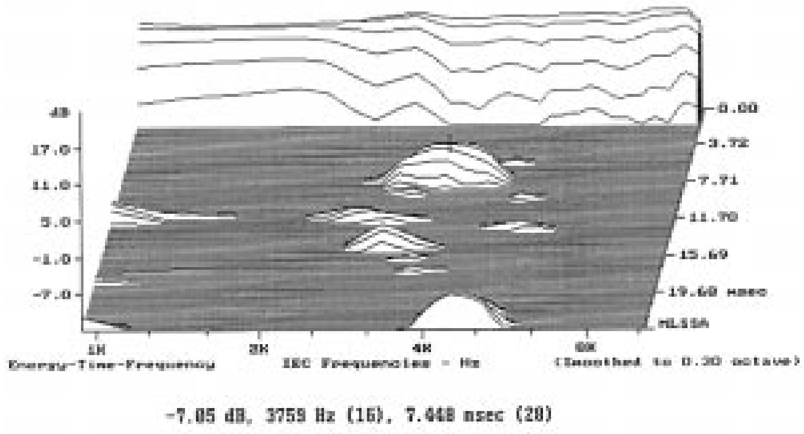

Fig. 3 - EFT response showing high level reflection at 7.5 ms and 3800 Hz.

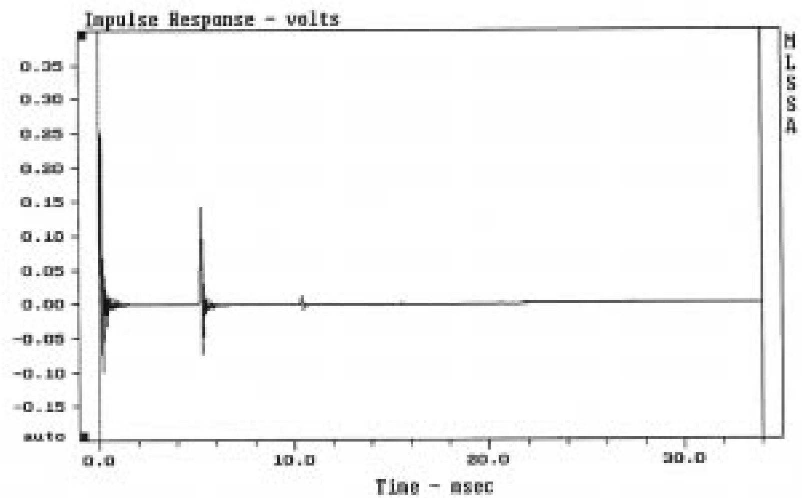

Fig. 4 - Impulse response of electronic reflection simulation.

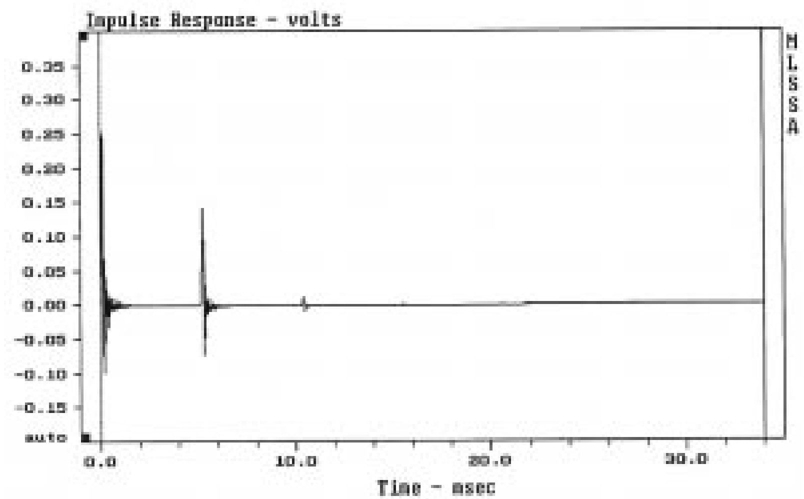

Fig. 5 - Fourier transforms of direct and delayed signal for electronic reflection simulation, 4.8 ms half-Hanning window.

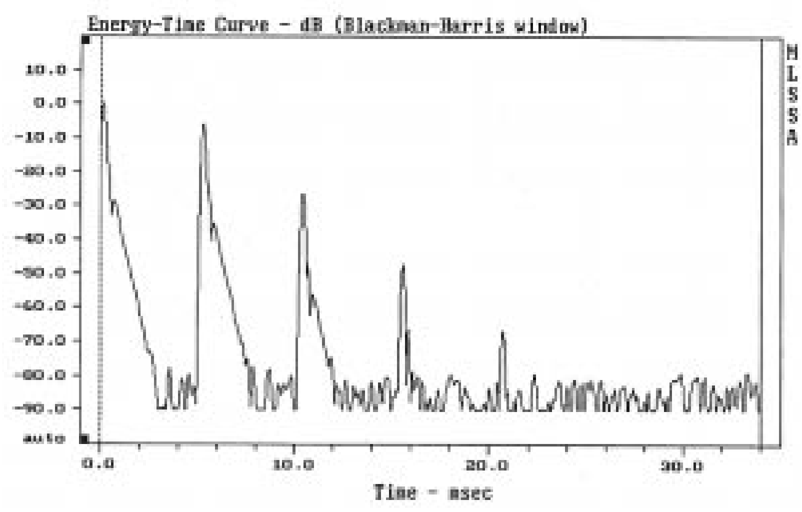

Fig. 6 - Wideband ET response of electronic reflection simulation.

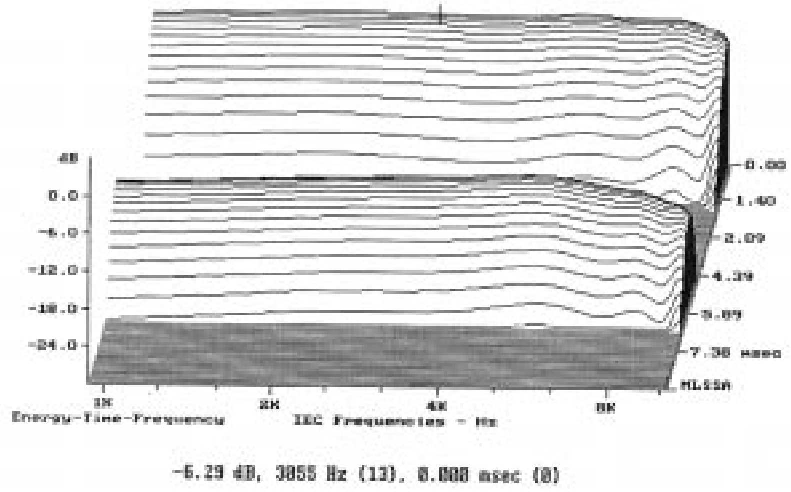

Fig. 7 - EFT response of electronic reflection simulation.

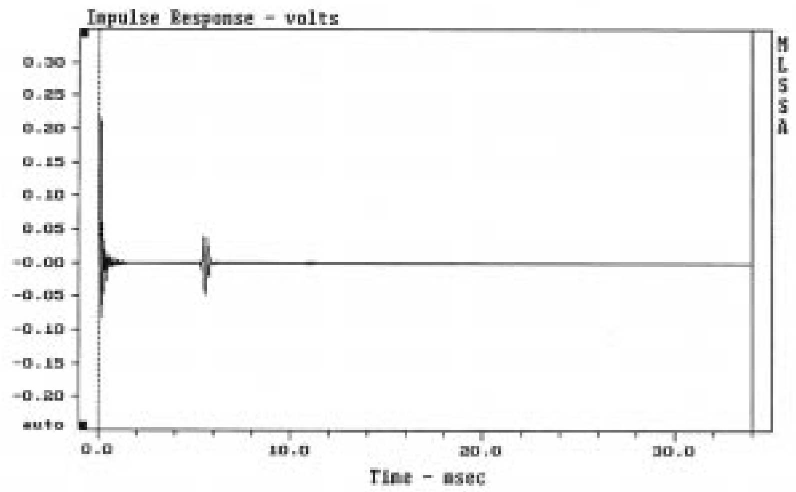

Fig. 8 - Impulse response of electronic reflection simulation with 4 kHz – 6 kHz bandpass filter.

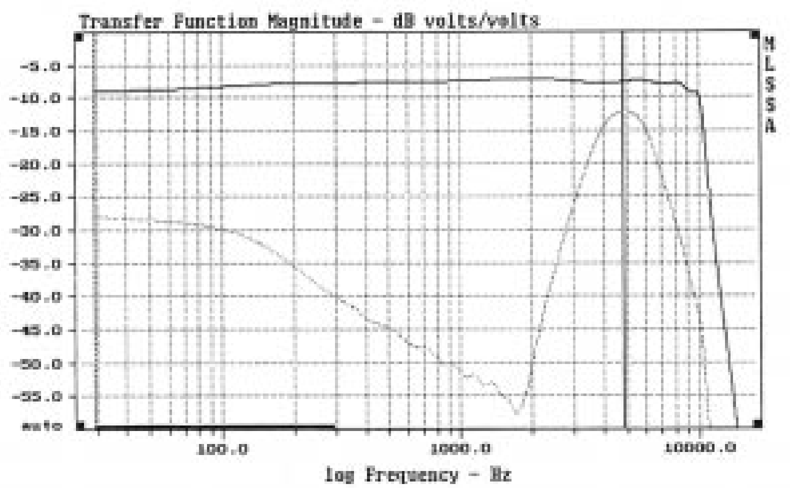

Fig. 9 - Fourier transforms of direct (upper) and delayed (lower) signals for filtered electronic reflection simulation.

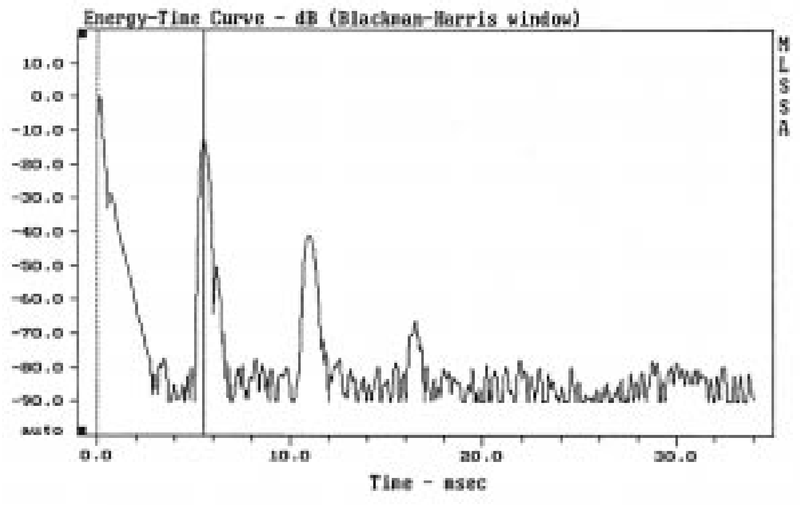

Fig. 10 - ET response for filtered electronic reflection simulation.

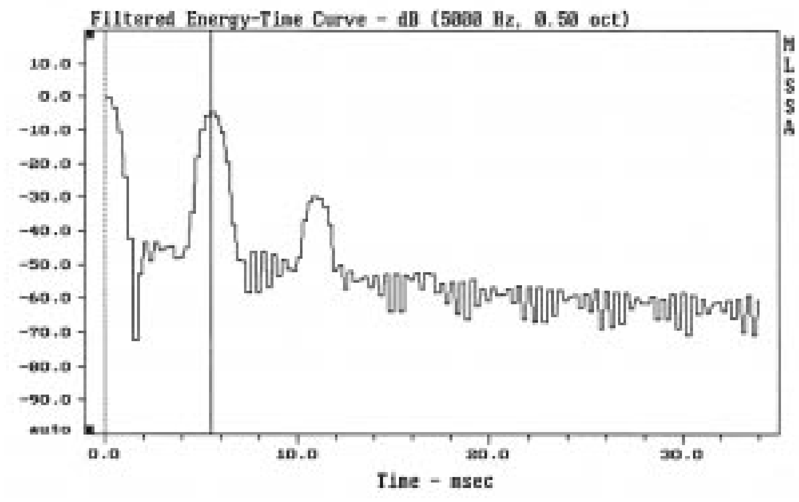

Fig. 11 - ET response, 0.5 octave 5 kHz filter, for filtered electronic reflection simulation.

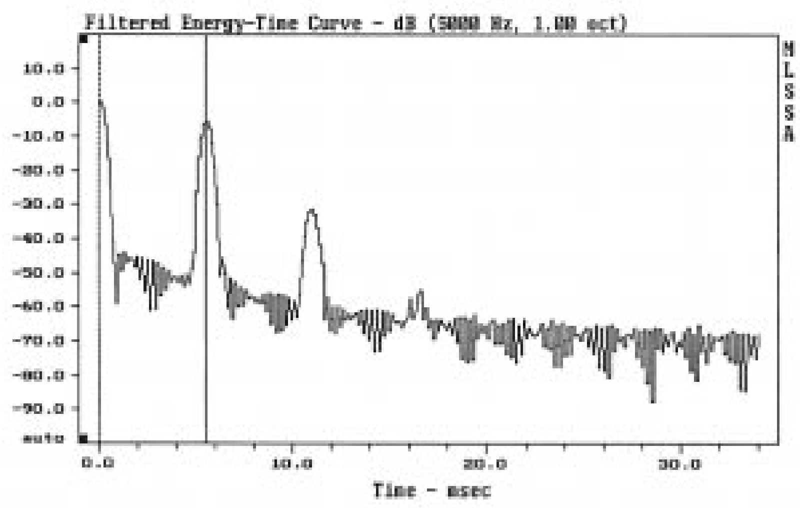

Fig. 12 - ET response, 1 octave 5 kHz filter, for filtered electronic reflection simulation.

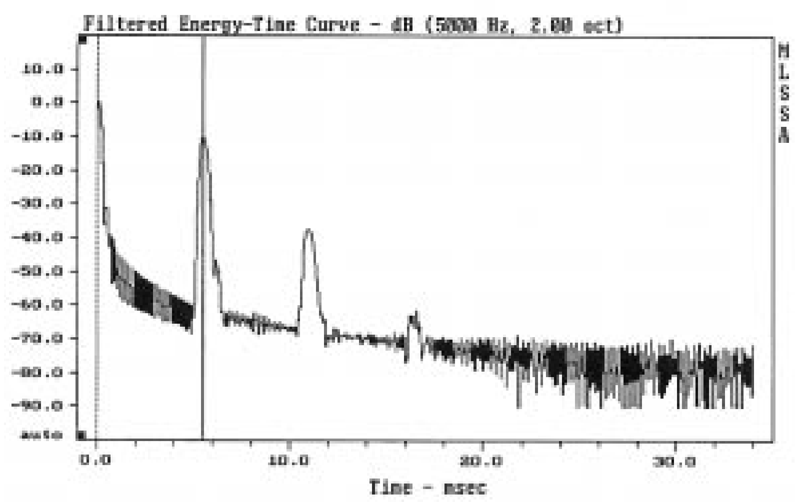

Fig. 13- ET response, 2 octave 5 kHz filter, for filtered electronic reflection simulation.

processor, at multiples of the delay length. These should be ignored for the present argument.)

6.2.1 Broadband delay

Fig. 4 *(see page 6)* shows the time response of the test system without the filter in the delay path. The 5 ms response is clearly visible. The later responses are also visible, up to the fourth order.

Fig. 5 *(see page 6)* shows the Fourier transforms of the main and the delayed signals. These were obtained using the maximum possible time window, just less than 5 ms, with rectangular weighting. The lowest reliable frequency, and frequency resolution, are indicated by the bar at the bottom of the graph (just over 200 Hz). The frequency response of the delayed signal is reasonably uniform. The ripples at low frequencies, up to 1 kHz, are an artifact of the windowing process. The delayed signal is between 7 and 5 dB below the direct signal.

Fig. 6 *(see page 7)* shows the ET response obtained using a Blackman-Harris window. The first peak, at 5 ms, is very close to −6 dB relative to the direct sound, as expected.

Fig. 7 *(see page 7)* shows the ETF response, obtained with a 4.26 ms half-Hanning window. Again, the delayed signal is 6 dB relative to the direct signal.

6.2.2 Band-limited delay

Fig. 8 *(see page 7)* shows the time response of the test system with the filter in the delay path. The 5 ms response is clearly visible.

Fig. 9 *(see page 7)* shows the Fourier transforms of the main and the delayed signals, obtained as before. The frequency response of the delayed signal can be seen to be centred on 5 kHz, with a relatively broad response, approximately −3 dB at 4 kHz and 6 kHz. (The apparent response below 1600 Hz is an effect of the windowing process.) The delayed signal is 4.6 dB below the direct signal at 4850 Hz.

Fig. 10 *(opposite)* shows the ET response obtained using a Blackman-Harris window. The first peak, at 5 ms, is apparently −12.8 dB relative to the direct sound. This is a measure of the average 'energy' levels of the two signals over the effective bandwidth of the measurement.

Fig. 11 *(opposite)* shows the ET response obtained using a half-octave wide filter centred on 5 kHz. In this case, the filter bandwidth, 4200 Hz to 5900 Hz, just encompasses the peak of the delayed signal response in the frequency domain. The apparent level difference of 4.49 dB corresponds closely with that obtained from Fig. 9. The significantly poorer time resolution of the narrow band filter is also evident.

Figs. 12 and 13 *(opposite)* show intermediate conditions, for a one-octave and a two-octave filter respectively. The progressive change in the ratio of the 5 ms reflection to the direct sound is clear.

Fig. 14 shows the ETF response, obtained as before. The delayed signal is apparently about −5 dB relative to the direct signal at the peak, approximately in accordance with the true condition. In this case, the effective bandwidth is less clear. It is probably about 338 Hz at ±3 dB (= ±1.44 Δf for a 128 sample record at 30 kHz sampling frequency). It is clearly narrow enough just to represent the passband component of the delayed signal at any one centre frequency. There is no inclusion or averaging of significantly remote frequencies.

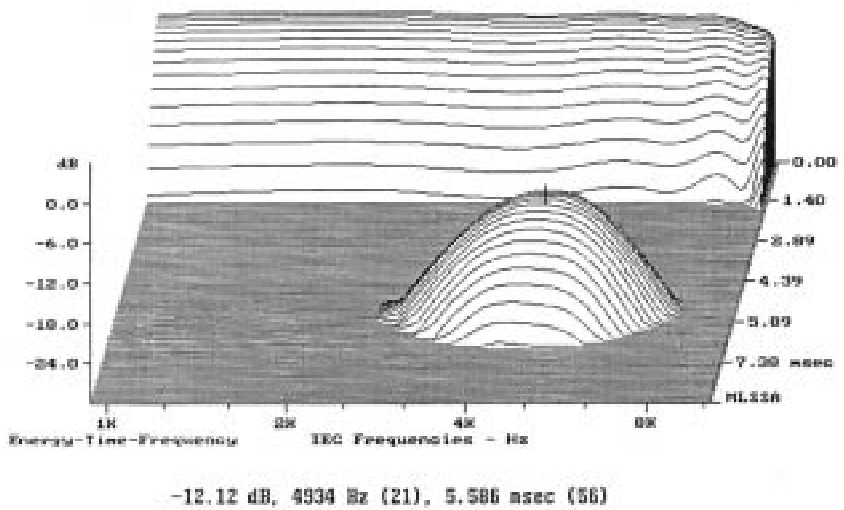

Fig. 14 - ETF response for filtered electronic reflection simulation.

7. DISCUSSION OF RESULTS

For any type of measurement it is self-evident that the result will be some form of average of all of the data which falls within the measurement scope, in time or in frequency or a complex combination of both. From the experimental model of a single room reflection, it has been shown that the measured results of ET and ETF responses were closely in accordance with the expectations, averaged over whatever frequency band was in effect.

Thus, if a measured reflection amplitude is obtained from either ET or ETF responses, the effective value obtained for the reflected- to direct-energy ratio depends on the bandwidth of the measurement and on the signal frequency responses.

The nominally unfiltered ET response will produce an overall value for the ratio of the two signals, averaged over essentially the whole frequency range. For a reflection (or a direct signal) which has pronounced frequency variations, the amplitudes of the highest components may be under-estimated. Narrower frequency bands may be used by effectively applying a frequency-domain bandpass filter before calculating the response.

The ETF response will always produce a relatively high frequency-domain resolution (at least in comparison with the unfiltered ET response), giving the true response at each frequency, within the limitations of the Fourier and Nyquist theories.

8. CONCLUSIONS

The fundamental resolution limits of time-frequency measurements have been described. It has been shown that the practical limits of time and frequency resolutions which can be obtained are adequate to quantify the early reflection patterns in small rooms. These limitations are close to, but probably just adequate to describe, those parameters necessary for a successful stereophonic illusion.

Experimental measurements, using an electronically simulated single room reflection, have demonstrated that the results obtained correspond with the theoretical expectations.

9. REFERENCES

1. WALKER, R., 1993. A new approach to the design of control room acoustics for stereophony. Paper G11, Preprint #3543, 94th AES Convention, Berlin, March.

2. SHORTER, D.E.L., 1946. Loudspeaker transient response. *BBC Quarterly*, **I***(3)*, October.

3. WALKER, R., 1979. A preliminary investigation into the measurement of time and frequency response of listening rooms and control cubicles. BBC Research Department Report No. 1979/9.

4. D'ANTONIO, P. *and* KONNERT, J., 1989. Complex time-response measurements using time-delay spectrometry. *J.A.E.S.*, **37***(9)*, September, pp. 674-690.

5. WALKER, R., 1995. Controlled Image Design: The management of stereophonic image quality. Research & Development Report No. RD 1995/4.

6. WALKER, R., 1995. Controlled Image Design: The results from the first installations. Research & Development Report No. RD 1995/5.

7. HARRIS, F.J., 1978. On the use of windows for harmonic analysis with the discrete Fourier Transform. Proc. I.E.E.E., **66***(1)*, January, pp. 51-84.

8. GADE, S. and HERLUFSEN, H., 1987. Use of weighting functions in DFT/FFT analysis (Part 1). *Bruel and Kjaer Technical Review*, **3**.

9. VANDERKOOY, J. *and* LIPSHITZ, S.P., 1990. Uses and abuses of the energy-time curve. J.A.E.S., **38***(11)*, November, p. 819

10. RIFE, D.D., 1989. Transfer-function measurement with maximal-length sequences. *J.A.E.S.*, **37***(6)*, June, pp. 419-444.

11. WALKER, R., 1990. Early reflections in studio control rooms: the results from the first controlled Image Design installations. 96th AES Convention, Amsterdam, Paper P12-6 (preprint no. 3853).

12. HEYSER, R.C., 1967. Acoustical measurements by time-delay spectrometry. *J.A.E.S.*, **15**, October, pp. 370-385.

APPENDIX

Swept tone versus impulse response measurements

A method of measuring system time-frequency responses was originally proposed by Richard C. Heyser[12]. This was based on the concept of excitation by a linear frequency sweep, and analysis by a filter whose centre frequency was offset from the instantaneous excitation frequency. Thus, the filter would respond (best) to a signal which was delayed in time by an amount related to the offset and the sweep rate.

In principle, the method is satisfactory. In practice, it suffers from significant drawbacks, even without comparison to other methods:

a) Theoretically, it is subject to the same time/frequency resolution limitations as any other method. However, the time-domain and the frequency-domain windows are not self evident. The time-domain window is clearly related to the impulse response of the offset filter, but not in a simple way, because it also depends on the frequency response of that filter.

b) It is time consuming — each different time offset requires a new scan. To obtain a full time-domain record requires the application of many scans. No post-processing is possible if additional data, for example a different time offset, is required.

c) Although many modern implementations use different processing methods (based on Fourier transforms of the cross-modulation product of the output and received signals, rather than implementing the offset filter directly), the windows are still not very self-evident, nor very well defined.

d) Most implementations, in practice, produce large quantities of spurious data which usually obscure the real data completely. This results from the poorly-defined windows.

e) The method produces the frequency-domain response directly. Any time-domain functions have to be obtained by inverse Fourier transformation.

In comparison, the direct measurement of the impulse response gives all of the information about the system in one simple result. The time-domain data is available directly. Fourier transformation has to be used to obtain frequency-domain data, but the windows can be readily applied and understood. The spurious data produced is within the theoretical limitations of the applied windows. Once the impulse response has been obtained, processing can be carried out without further measurements. By recording just the time-domain response data, all of the necessary information can be saved, for additional post-processing at any time in the future.

The direct measurement of impulse response carries a significant signal-noise ratio penalty. However, the use of maximal-length sequences and cross-correlation, as employed in MLSSA, makes the signal-noise ratio performances of swept-tone and MLSSA about equal.

Published by BBC RESEARCH & DEVELOPMENT DEPARTMENT, Kingswood Warren, Tadworth, Surrey, KT20 6NP