Non-echoic acoustic measurement with the H-P 3582A

New Hewlett-Packard spectrum analyser uses digital signal processing

by **R. N. Grubb,** Auris of Boulder, Colorado

The HP3582A is a recently announced audio spectrum analyser using fast Fourier transform analysis. A number of its features can be exploited in the measurement of loudspeakers and microphones in non-echoic conditions. These are described and some practical examples of its application given.

THE RECENTLY announced model 3582A spectrum analyser by Hewlett-Packard is an example of the new generation of instruments which depend on microprocessor technology to provide powerful capabilities at a lower price than has previously been possible. In this case, digital signal processing technology is used to implement a flexible 0.02 Hz-25.5kHz spectrum analyser, using the fast Fourier transform (FFT) of the digitized input signal to calculate the signal spectrum in the frequency domain from a sample of the input signal in the time domain. Although the instrument is a computer system, the mechanics usually associated with the use of a computer are completely transparent to the user, who is presented with a fairly conventional-looking front-panel control layout. The program is, of course, contained in read-only memory.

The 3582A is not a real time third octave analyser. In fact, one thing which may put off the average audio engineer is the lack of anything but linear frequency-scale presentations. However, it is inherent in the fast Fourier transform approach that a linear, equally-spaced set of spectral estimates is produced. The resolution and bandwidth of each estimate depends on the length and shape of the time window used to select the signal sample for analysis. Thus a logarithmic presentation of the data would necessarily be only cosmetic, information at the higher frequencies being lost, if a constant proportional resolution were displayed. As the available frequency ranges of the instrument are very extensive, all the information is available, although it is perhaps more time consuming to obtain.

By audio spectrum-analysis standards, the capabilities are unconventional, including measurement of phase, measurement of transfer functions and time-domain signal averaging before analysis.

Measurement of the phase response of audio systems, particularly of loudspeakers, has recently become of interest in the quest for the more realistic reproduction of transients. The 3582A provides in one box the means to make response measurements, including phase, on loudspeakers and other audio trandsucers, without requiring an anechoic chamber, or the roomful of minicomputer used by loudspeaker manufacturers to make such measurements.

Fig. 1. Arrangement required for phase measurement.

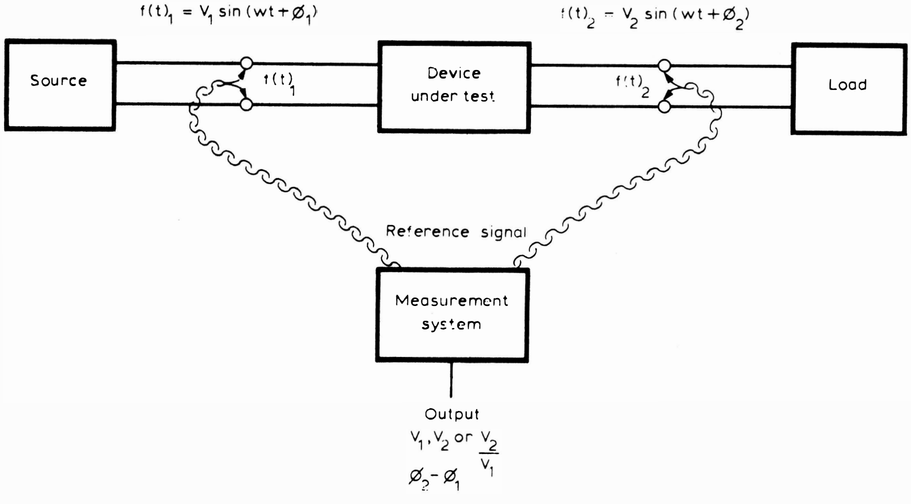

Phase

Before proceeding to explain how to use the analyser for this hurpose, it may be useful to some readers to review what is meant by phase response and in particular how it can be measured by a spectrum analyser. The phase response of a device refers to a measurement of relative phase, usually the difference between the input and the output of the device. Unlike amplitude, or spectral amplitude, which is measured with a single connexion to the system under test, two separate connexions are needed to measure phase response as in Fig 1. Thus, although a spectrum analyser is normally a single-input device, with analysers like the 3582A, one must think in terms of two inputs to measure phase. Simply feeding in a composite signal to one channel of the instrument will give a perfectly good amplitude spectrum, but the phase answer computed will be different for each time sequence analysed because of the lack of a reference. This may not matter in some applications. For instance, if we want to know whether sidebands observed on a carrier are due to amplitude or phase modulation, their phase relationships to the carrier itself as seen in Fig 2 and a single sample

Fig. 2. Identifying amplitude or phase modulation.

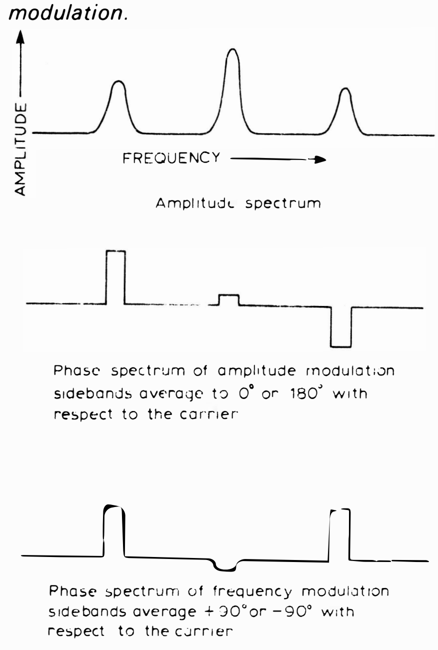

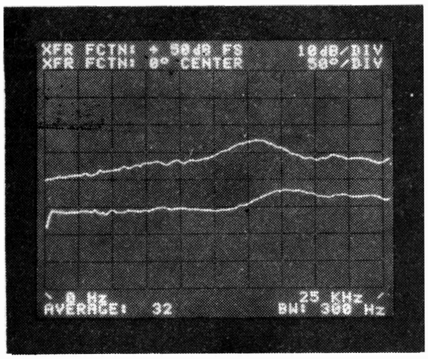

Fig. 3. Comparison of two AKG C451 microphones. Lower trace-amplitude. Upper trace-phase.

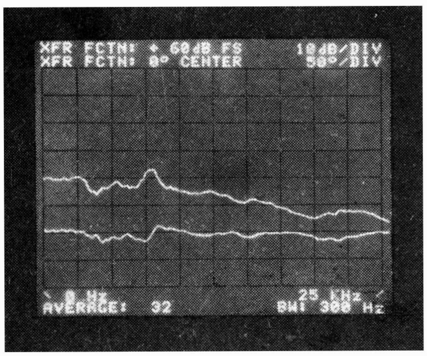

Fig. 4. Comparison of two 1in diameter capacitor microphone capsules in a stereo coincident pair configuration. Lower trace-amplitude. Upper trace-phase.

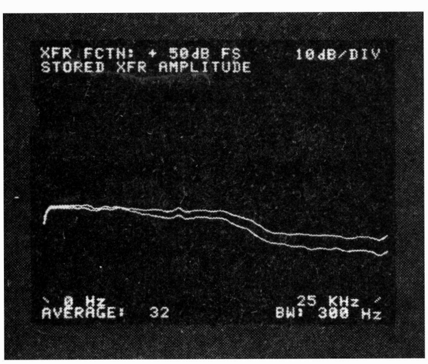

Fig. 5. The effect of a foam windshield on a C451. Upper trace — without windshield. Lower trace — with windshield. (Scales are the same as Figures 3 and 4).

analysis will give us the answer we want.

Transfer function

The most straightforward mode of operation to give repeatable phase measurements is that of the transfer function measurement. The two channels of the analyser are connected across the input and the output of the device to be tested and one of the two built-in noise sources connected to the input. The analyser now plots the ratio of the amplitudes and the difference in the phase of its two inputs versus frequency.

This transfer function measurement capability can be applied very neatly to the measurement of microphones. By connecting two microphones, one of which is to be regarded as the standard, to the two inputs of the analyser and placing them close together and in the sound field of a loudspeaker fed from the analyser noise source, their responses can be compared directly and very quickly. Figure 3 shows the result of comparing two nominally-identical C451 microphones with CK1 capsules. This disclosed the interesting information that although the microphones are well matched up to 15kHz, the two differed by nearly 6 dB at 17.5KHz. In this case, since neither microphone could be regarded as a standard, it was not possible to say which microphone or whether one or both was at fault. Exchanging the capsules on the microphone bodies showed the problem to be in the capsules and not in the microphone electronics or the amplifier chains.

The upper trace shows the phase difference. The constant phase slope at low frequencies shows that the "test" microphone was slightly in front of the reference microphone and it was possible by careful adjustment of the relative microphone position to make the phase slope zero. It is interesting that the difference between the capsules shows up in the phase at a lower frequency than in the amplitude. One thing to note in this and in most of the other examples shown is that the lowest-frequency point plotted by the analyser in the zero frequency start mode is in fact actually 0Hz, i.e., d.c. and the position of this point depends on the analyser amplifier d.c. offsets or externally applied d.c. In this case, of course, the microphone amplifiers are a.c.-coupled, so the zero frequency point is quite meaningless.

Figure 4 shows another comparison of two microphones, in this case two 1in capacitor capsules mounted one above the other in the same case and designed to be used as a coincident stereo pair. The lower trace is the magnitude again. This showed a good match at all frequencies, except in the region 3-9kHz, where there are 2-3dB differences. Some experiment and the use of another microphone as a comparison standard showed that the irregularities were only present in the lower of the two capsules and were very sensitive to the angle of the microphone in the vertical plane to the direction of the incident sound field. This seemed to show that the problem was due to diffraction effects at the microphone case, the lower capsule being much closer to the case than the upper.

Yet another interesting comparison is shown in Fig. 5. This is the pair of C451s again, but this time the stored trace facility has been used to show the effect of the standard foam windshield on one of the microphones. The effect is easily measurable and amounts to nearly 3dB at 15kHz (unfortunately, I forgot to illuminate the graticule for this photograph!)

Impulse testing

All the preceding three examples were measured in a normal room with some acoustic treatment, but nevertheless far from anechoic. Thus, the sound field at the microphones being compared is composed of direct and reflected components. The comparison results have to be based on the assumption that the microphone polar responses are similar. It is only possible by this method to compare a cardioid microphone with another cardioid or an omnidirectional one with another omnidirectional microphone, etc. Providing the pair of microphones is not too far from the source compared with the dimensions of the room, and that the room is reasonably non reverberant, then small errors in polar response should have little effect on the comparison. However, we can do this kind of measurement in a non-anechoic room without these restrictions by using the capability of the instrument to analyse the impulse response of loudspeakers and microphones and present the results in the more familiar terms of amplitude and phase and it is to this, probably least familiar, mode that I now turn.

Fourier transform theory tells us that a zero width pulse contains equal energy per unit bandwidth (power spectral density — p.s.d.) at all frequencies, i.e., it possesses an infinite bandwith. Of course, this is a mathematical abstraction because, unless the impluse is infinitely large in amplitude its energy in any particular bandwidth will be infinitely small. Fortunately for any given audio bandwidth, it is easy to produce an impulse sufficiently narrow for the p.s.d. to be flat. The theory tells us that the power spectrum of a pulse of width t is

$$P(f) = \left(\frac{A\sin\pi ft}{\pi ft}\right)$$

This function, the familiar sin x/x, is plotted in Fig. 6. By choosing t to be small enough, we can make the p.s.d. as flat as we wish over the working bandwidth. For instance, it is easy to calculate that a 1μs wide pulse is only 0.01dB down at 25 kHz, the maximum band-

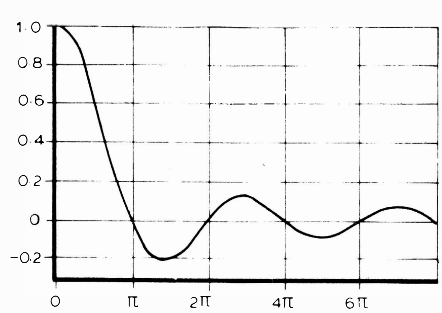

Fig. 6. The function sin x/x.

width of the analyser. A 10μs pulse is only ≈ 1dB down. At the rear of the 3582A is a t.t.l.-level impulse output. This gives a positive-going pulse which is≈ 1μs long at the widest analysis bandwidth (25kHz) and which increases in width as the analysis bandwidth is reduced. If this output is connected to the input of the analyser, the displayed amplitude spectrum will show the first of the problems of impulse analysis which has to be carefully considered in order to obtain valid results. Indeed, the analyser shows a flat spectrum but, as the sensitivity is increased to bring the observed spectrum above the baseline the input channel overload light rapidly comes on. In fact, it is impossible to get more than a 20dB measurement range above the noise floor. This, of course, is because the test signal has a very high ratio of peak to mean value, and the analyser input dynamic range, which is set by its analogue to digital converter, only permits this limited range in the spectral domain. This situation can be improved considerably, however, if an external impulse source is used. As calculated above, a pulse of ten times the width (10μs) is about 1dB down at 25kHz. This gives another 20dB of analysis dynamic range, which is adequate for nearly all acoustic testing; it is easy to correct for the small loss at high frequencies of the test signal, if 1dB is important.

Phase

Having developed the test signal, the next question to consider is what is meant by the phase of the test signal and how the analyser measures it. The reference, in this case, is set in the time domain by the position of the time window, in which the analyser samples the input signal. At a time t_0 one can think of all the reference frequencies starting simultaneously at zero phase (zero amplitude for a cosine wave). If the impulse is positioned at t_0, then its spectrum consists of all frequencies also starting at zero phase and the analyser will read 0° at all frequencies. If the impulse is displaced from t_0 then there will be a progressive displacement, increasing with frequency, in the analysed phase expressed by the formula for the group delay introduced by the displacement

$$\frac{\Delta\phi}{\Delta f} = \Delta t \times 360°$$

(ϕ in degrees, f in Hz.)

where $\Delta\phi/\Delta f$ is the phase slope with frequency. For a positive delay (signal later than t_0) the phase of the higher frequencies lags the lower and vice versa. Note that a linear rate of change of phase implies only a delay and no waveform distortion.

In the 3582A, t_0 is set at the middle of the time window when the 'flat top' or Hanning passband shape is selected, or at the start of the time window when the 'uniform' passband shape is selected. The latter is the passband intended for transient analysis. In the former cases, the passband shape is set by amplitude weighting in the time domain so that a transient at the beginning or end of the time window would not be analysed correctly. To be able to interpret the phase readout from the analyser, it is necessary to place the impulse close to t_0, because a large phase slope due to a time difference will obscure the properties of the system under test and, if too large, renders it discontinuous, because the discrete samples computed by the analyser are not close enough together to resolve the rapid phase change. To adjust the timing, the analyser can be operated in two ways and can be thought of more like an oscilloscope. In fact, the time-domain sampled waveform can be selected for display on the c.r.t.; this is an almost indispensable mode for setting up the analyser for transient analysis. In the free-run mode, the instrument repeatedly starts new time windows as soon as it is ready to analyse new data. The rear-panel impulse outout occurs at the start of each time window. Alternatively, the analyser can be triggered like an oscilloscope by an input signal on channel A or by a t.t.l. level pulse at a rear-panel input.

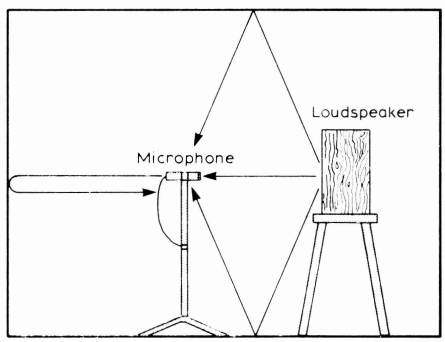

Fig. 7. Sound paths for direct and reflected sound in a small room.

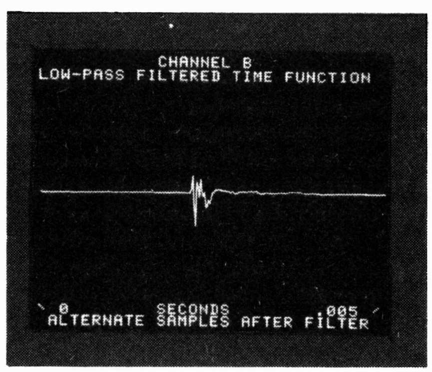

Fig. 8. Typical loudspeaker time domain response when driven by an impulse.

Echo gating

The advantage of using a transient signal to analyse the response of acoustic devices is that it is possible to suppress the effect of room reflections entirely without having to work in an anechoic room. To a close approximation, sound travels 1 foot per millisecond: the typical response of a loudspeaker to a 10μs wide impulse is over in 2-3ms, depending on the physical size of the cabinet. Even in a quite small room with a loudspeaker 3 to 4 feet from the floor and the measuring microphone 8 feet away, the first room reflection will arrive at the microphone 3-4ms later than the direct sound. Figure 7 shows the situation. A typical time domain response of a loudspeaker to a 10μs wide impulse is shown in Fig 8, which was taken from the analyser screen, with the instrument set on the 0-25kHz range. On this range the time window is ≈ 5ms long and, by controlling the trigger time, the transient picked up by the measuring microphone can be positioned near the centre of the time window with the first reflection just outside the window. This enables the amplitude response to be obtained, but as explained above, the transient should really be positioned near the start of the time window if the phase response is desired. Since the time window gets longer as the analysis bandwidth is reduced (necessary if the

Fig. 9. Timing diagram for impulse measurements on loudspeakers.

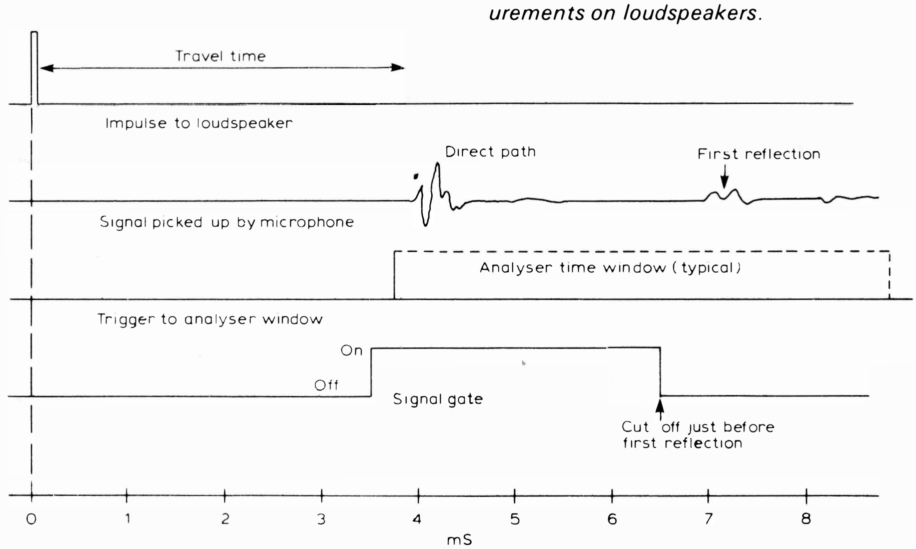

low frequency response is to be examined in detail), an electronic signal gate is needed so that the first direct-path signal can be isolated. To do this, and to be able to adjust all the delays correctly and generate the test impulse required some auxilliary equipment in addition to the analyser itself. This is unfortunate because it seems that it would have been quite simple to build all the required functions into the analyser in the first instance.*

Figure 9 shows the overall timing and gating required. Because the analyser time window must be started later than the impulse sent to the loudspeaker, it is best to generate the measurement repetition rate externally. This should be set to the highest rate which allows all room responses to die out before the next pulse.

Two delayed trigger pulses are then needed – one to start the analyser time gate at the correct time with respect to the transient picked up by the measurement microphone, and one to start the signal gate. A convenient way to get the first delay is to use a second microphone slightly closer to the loudspeaker under test and feed its amplified output to channel A of the analyser as the trigger signal. The measurement microphone output is fed to channel B. The delay is adjusted by setting the relative distances of the two microphones to bring the received transient just at the start of the time window on channel B. Channel A should also be examined to make sure that the trigger point on the transient is a stable one.

It is very important to make sure that all the significant energy from the transient radiated by the loudspeaker is included in the time gate. This can be checked both by inspection in the time domain and by changing the signal gate window over a small range and seeing if it affects the transformed frequency and phase response. With high quality loudspeakers of small dimensions, it seemed the response died essentially to zero after about 3ms, and it seemed to be possible to get a clean separation between the direct arrival and the first reflected arrival in a room with a smallest dimension of 8 feet. With larger loudspeakers or units with pronounced resonances, this may not be possible and it would be necessary to use a larger room.

The delay mechanism for the signal gate and the signal gate itself need to be electronic Commercial pulse generators can be used to generate these and the basic impulse and its repetition rate or, with the aid of a few digital i.cs, a special generator and controller could be assembled. Some commercial signal gating devices may be satisfactory in this application – a simple shunt f.e.t. switch such as is shown in Fig. 10 works well. It is most important that the switch does not introduce appreciable transients itself in the signal path. When the analyser bandwidth is reduced, the time window becomes longer and it may be necessary to readjust the system repetition rate. Also, as discussed previously, the impulse length must be increased proportionately to preserve approximately constant spectral power density.

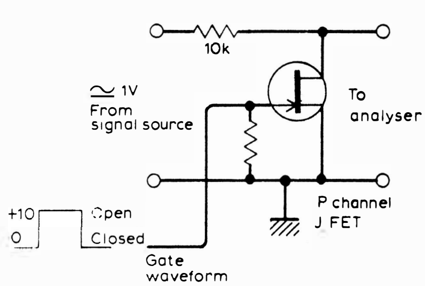

Fig. 10. F.e.t. signal gate.

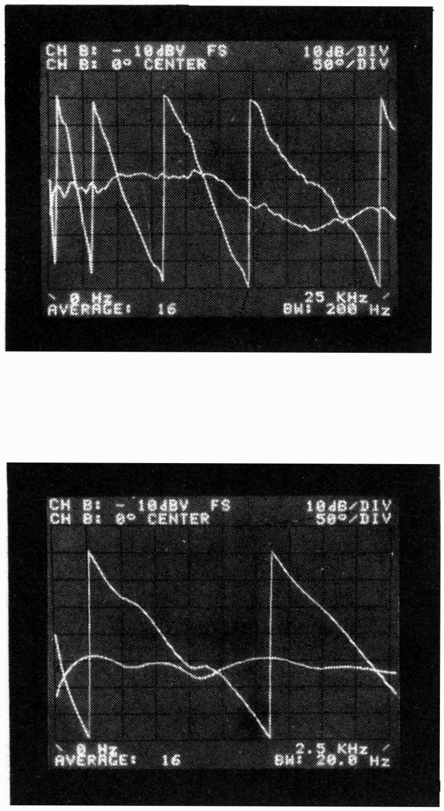

Fig. 11. Frequency and phase response of a Spendor BC1 loudspeaker measured with an impulse, a) 0-25kHz, b) 0-2.5kHz.

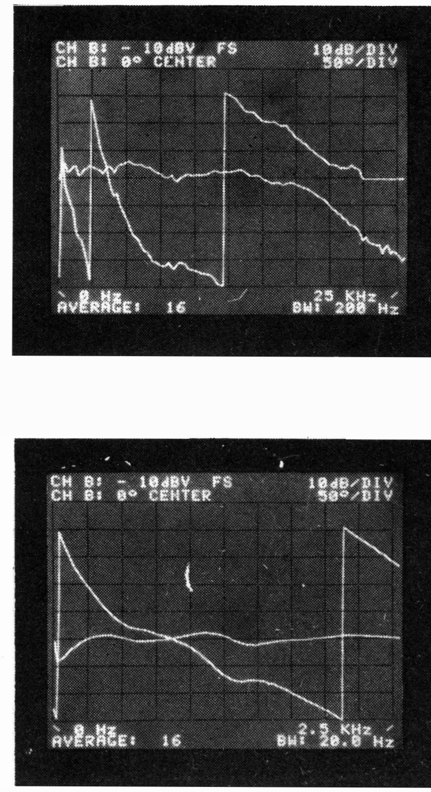

Fig. 12. Frequency and phase response of a Chartwell LS3/5A loudspeaker measured with an impulse. a) 0-25kHz, b) 0-2.5kHz.

Practice

Unfortunately, no measuring technique is completely free of disadvantages and the gating-out of room reflections is no exception. The problem is that of determining whether the initial response of the loudspeaker really has died away or not. It turns out that the use of a time sample of length t produces an uncertainty in the value of the spectral amplitude points for all frequencies roughly less than $1/t$ in frequency. Why the effect is an uncertainty and not just a calculable loss can be seen by considering a couple of simple examples. If the device being analysed is perfect (i.e. a piece of wire) then locating the time window would clearly have no effect, because the input impulse signal has a zero value at all times except for a small interval near zero time. However, if the device had a low-frequency cut off caused by the equivalent of a single pole RC network, then its response to the impulse would have an overshoot following the impulse which returns to the baseline exponentially with a time constant of RC seconds. In this case, a significant error will be made in the low-frequency response measurement unless the time window is maintained for 5 or 6 time constants, so that the response has reached zero for all practical purposes. Locating the impulse response at a point where the net remaining area under the response is negative will result in an apparent enhancement of low frequencies well below $1/t$ and vice versa. Thus the effect of the truncation depends entirely on the exact form of the impulse response.

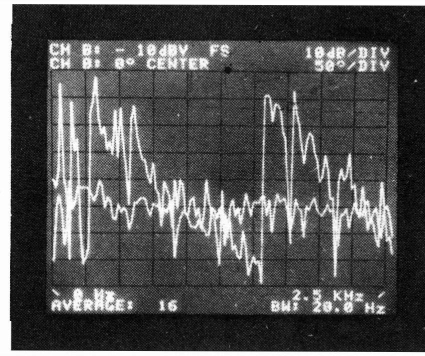

Fig. 13. Apparent response of a BC1 with the signal gating disabled and the first room reflection included.

Figures 11 and 12 show typical results obtained in the author's studio with a Spendor BC1 and a Chartwell LS3/5A. The phase responses clearly show the effects of the crossover in the case of the Chartwell as a change in time delay (phase slope) starting at ≈ 3kHz. With a little more flexible arrangement, the average phase slope could have been brought closer to zero with resultant ease in interpretation. In all cases, the measurement microphone was about on axis and 6 feet from the loudspeaker. Figure 13 shows the effect of disabling the signal gate and allowing some of the room reflection to be analysed! I found that these loudspeaker measurements were relatively unaffected by the microphone used, providing it was a capacitor type and of professional quality, since these microphones invariably have a much flatter response than monitor loudspeakers. If a standard measuring microphone is not available, then a ½in diameter, omnidirectional capacitor microphone such as the AKG C451 with a CK2 capsule would be the best second choice. The examples shown were made with this same microphone but with a CKI capsule, which probably does affect the results somewhat. In all cases the low frequency response below about 2-300 Hz appears to be attenuated compared with the published responses of these particular speakers, so it must be assumed that some truncation of the impulse response was taking place.

Care should be taken not to overdrive the loudspeaker with the impulse: A few watts peak power should be all that is required. The sound should be that of a quite quiet tick similar in volume to that of a typical alarm clock. If the measurement conditions are quiet, then the response can be obtained with only one impulse. However, if you don't live in the country or have a well isolated studio handy, there is no need to despair; use the last unique feature of the analyser, time domain averaging. This adds together, algebraically, each successive sample at the same time with respect to the trigger. The wanted signal is preserved but non coherent background noise cancels itself on the average. Thus, not only do you not need an anechoic room to make loudspeaker measurements, you do not even need a quiet one. The examples in Figs. 11 and 12 used a signal average of 16 impulses.

All the comparison tests of microphones described earlier can be better done using a loudspeaker excited by an impulse with the appropriate delays and gating. In this case, since both signal channels will be needed for the measurement, the rear-panel t.t.l.-level trigger input must be used. Absolute measurement of microphone response requires an acoustic impulse generator of known characteristics. It has been reported in the literature that a high-voltage spark discharge or an exploding wire forms a useful source for this purpose, providing the construction of the electrodes is such that the sound radiation is unimpeded. However, the author has not yet tried this.

* It is possible to do the signal gates within the instrument using the IEEE 488 bus programing input. However, this means significant additional complication and expense.

References

[1] Hewlett Packard Application Note. Understanding the HP3582A Spectrum Analyzer.

[2] J. M. Berman and L. R. Fincham. The Application of Digital Techniques to the Measurement of Loudspeakers. *Journal of the Audio Engineering Society*. June 1977, Vol 25, No. 6, pp. 370, 384. □

Abridged specification

Input channels

2n 1MΩ + 60pF impedance, sensitivity +30dBV to −50dBV in 10dB steps. Overload indicator light.

Frequency spans

1 Hz to 25 KHz full scale in zero-start mode, 1-2.5-5-10 sequence. Bandpass mode 5Hz-25kHz span in 1-2.5-5-10 sequence.

Frequency resolution

256 spectral points are calculated in the single-channel mode, 128 in the dual-channel mode. The resolution depends on the passband selected. There are available, a "Flat top" optimised for harmonic analysis of tone signals, a "Hanning" passband optimised for general random noise measurement and a "Uniform" passband intended for transient analysis and use with the built-in periodic noise source.

Display

The digitally driven c.r.t. has infinite storage capability. It can display up to two traces of data from either the current measurement or from up to two traces stored from previous measurements. It provides an alphanumeric readout of trace calibration, a cursor readout of trace values in engineering units and error messages. Amplitude display 10dB or 2dB per division (8 divisions vertically) or linear. Phase ± 200° Frequency displayed linearly.

Measurement modes

1. Frequency spectrum, amplitude and phase
2. Transfer function, ratio of input channel amplitudes and difference in phase.
3. Coherence function, the degree of coherence (0-1) between the input channels.

Signal sources

1. Random noise. This is generated digitally and adjusts automatically with the frequency range selected to maintain a constant power output in the analysis band.

2. Periodic noise. This is also generated digitally and is arranged to have a "comb" spectrum which exactly matches the calculated spectral points. This gives the same effect as a tracking generator in a conventional swept analyser with the advantage over random noise that no frequency domain averaging is needed to get an accurate answer.

3. Impulse. This varies in width depending on the frequency range selectedm It is timed to occur at the start of each analysis time window.

Averaging modes

1. Frequency Domain

a) r.m.s. average of calculated spectral points with 4-256 points averaged or an exponential "running average" mode.

b) Peak, 4-256 points or peak hold in a continuous mode.

2. Time domain.

4-256 input signal time sequences averaged. The zero time is set by a trigger circuit on input channel A or by an external trigger input at t.t.l. level.

EEV provides bright lights for ATV games

A large scale, computer-controlled electronic display board supplied by English Electric Valve Co. can be seen by television viewers of the Bob Monkhouse "Family Fortunes" panel game on Sunday evenings.

The main body of the display consists of 300 "character display tubes" (a form of c.r.t. costing about £100 each), which EEV say offer very high variable brightness, low power consumption and electronic switching with low level logic. The control logic, including a keyboard and v.d.u. control console, includes an Intel single board computer and the complete installation is said to have cost ATV about £80,000.

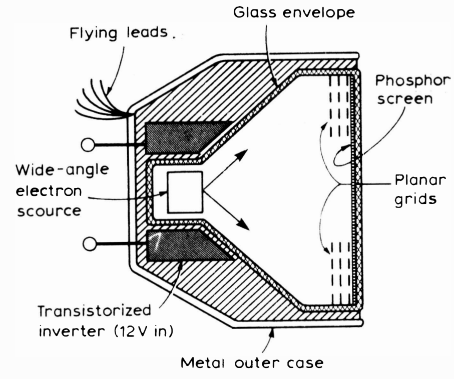

Cross-sectional view of the EEV character display tube. The flying lead grid connections are for multiplexing; the "expected life" of the tube is 40,000 hours or about five years.