Proceedings of the Institute of Acoustics

APPROXIMATION FUNCTIONS FOR VIRTUAL ACOUSTIC MODELLING.

R. Walker BBC R&D, Kingswood Warren, Tadworth, Surrey, UK. KT20 6NP.

1. INTRODUCTION.

At the present time, virtual production is becoming commonplace in television – for the video. For a convincing illusion, it is also necessary to model the acoustic aspects of the virtual space but, up to the present time, the audio has not received the same development effort as the video. Although many research groups are working on various forms of auralisation, virtual audio production for television has particular requirements, especially in terms of the speed of response, to allow tracking of moving objects.

The objectives for Virtual Production audio are simpler than many other auralisation paradigms in that the model does not have to represent a real room exactly. It merely has to present the audience with a credible acoustic illusion. That is helped greatly by the presence of the visual cues. A previous paper [1] described the principles of such acoustic modelling and the development of an experimental virtual audio system for broadcasting. The system included the modelling of some acoustic effects using simple filter topologies. It also discussed the need for simple implementations in order to achieve high rates of scene update and described some of the types of acoustic propagation effects that should be modelled.

The simple model included contributions for the direct sound, six discrete early reflections (representing the first-order reflections from the boundary surfaces of the room) and either obstruction or reflection of the direct sound by an internal object within the room. It also included contributions for the overall, diffuse reverberation. Provision was also made for directional source radiation characteristics.

This present paper describes, in more detail, some of the filter approximations used and their derivation. It especially discusses the principles of simple approximations to the effects of diffraction around obstructing objects. Other approximations described include air absorption as a function of distance and human voice directionality.

2. ACOUSTIC PROPAGATION EFFECTS.

A simple geometrical acoustic model, as described, would produce delayed versions of the input signal representing the discrete direct sound and early reflections together with some level of diffuse, reverberant sound. As a first step towards acoustical realism, the signals must be filtered to model a number of acoustic propagation effects.

In the simple model, all of the filters were implemented as a combination of a wide-band attenuation and a first-order, IIR low-pass with a single sample period of delay. Fig. 1 shows the basic filter arrangement, where α is the overall gain constant and β is the frequency response coefficient.

2.1 Air Absorption and Wall Reflection.

As sound propagates through air, the higher-frequency components are selectively absorbed, mainly by a complex, oxygen-water vapour resonance. Under common atmospheric conditions and at ordinary acoustic frequencies, this produces a progressive (monotonic) high-frequency loss.

As sound is reflected from a surface, a further selective loss of certain frequencies will occur if the reflection coefficient of the surface has a frequency dependant component. In principle, that modification of the frequency response might be arbitrarily complicated. For the simple model, only high-frequency loss was implemented.

Filters to implement that progressive loss due either to air absorption or selective surface absorption were included in the early sound processor. Each of the seven discrete sound components had a separate filter to model different distances and types of surface.

2.2 Source Directivity.

Most real sources have non-uniform directional radiation characteristics. The range of potential source types is very large. However, the majority of applications of Virtual Production will involve human speech. The only directional source included in the simple model was an average human speaker.

Source directivity was implemented by including a low-pass filter corresponding to the radiation angle from the source in the appropriate direction. That might be either the direct sound or one of the wall reflections. Therefore, each of the seven component signals of the early sound had a second, separate filter (in addition to that for air absorption).

2.3 Reverberation.

The basic reverberation processor produced a response based on the size of the room and the average, wideband absorption coefficient of the room surfaces. The processor included two frequency-response modifying stages. The first was an input filter to set the bandwidth. The signal passed through that filter only once so it could be used to implement features depending on a single-pass, such as the spectrum of the source (it was actually used to implement a frequency-dependant directivity index for source directionality).

The second was a pair of filters inside the signal recirculation loops. They could be used to implement those factors that progressively increased with time/distance, such as air absorption.

By effectively either under- or over-compensating for air absorption, a limited range of different types of reverberation frequency characteristics could be produced.

2.4 Internal Objects.

The basic acoustic model represented only the interior surfaces of the empty room. The representation of a multiplicity of objects within the room would have been relatively simple in principle, but the complexity of the model would very rapidly have become unmanageable. However, one of the most striking subjective effects in the acoustics of a room is the effect of a source being obstructed. In order to allow for sound sources passing behind or close in front of other objects (as required for flexibility and realism in production), it was necessary to model at least one internal object.

Such an object could either obstruct the direct sound or add an additional reflection to the six from the boundary surfaces. No attempt was made to model second-order effects, such as obstruction or re-reflection of a boundary reflection. The effects of obstruction in particular on the spectral response of an acoustic path can be very pronounced. The subjective effects are further emphasised by the large reduction in the direct path contribution.

In the remainder of the model, no account was taken of diffraction, even though in reality it has a very profound effect on the way sound behaves physically. However, in the modelling of obstructing or reflecting objects, especially the former, the effects of diffraction are what control the basic response. A listener or source entering a shadow zone behind an obstruction perceives a progressive loss of high-frequency sound, as well as some overall (wideband) attenuation. If that were not modelled, at least to some extent, then the subjective result would be unconvincing. It is especially so because the main signal (by definition, the direct sound) would be significantly reduced in level.

For obstruction, defined as when the physical object came close enough to the line from the source to the receiver to have a significant effect (not the same thing as geometrical obstruction), the direct sound image was effectively switched off and replaced by two additional images of the source. Those images were located at the correct time delay, corresponding to the total path length from source to receiver around each edge of the object, and at the geometrical angles subtended at the listener by the two edges of the object.

Actually, the replacement of the direct sound with two diffracted sound rays took place in two stages. As the line between a source and a receiver approaches obstruction, the effects of diffraction cause the frequency response and overall sound level to change before the point of actual geometrical obstruction. In that zone (the Fresnel zone), the source direction and distance remain those of the real source but the signal has to be filtered and attenuated. Beyond the point of geometrical obstruction the direction and distance become dependant on the relative position of the object edge (and the filtering increases). These partial and full obstruction calculations and source image substitutions had to be carried out for both edges of the object, complicating the detection of the obstruction condition and selection of the appropriate filters.

For reflections, the single image location was calculated and one of the 'obstacle' channels used to add another image at the appropriate distance and angle. In that case, the direct sound remained active. For reflections, the requirement for frequency-selective attenuation was less significant, mainly because, by definition, the direct sound remained as the dominant component. Thus, the contribution from a reflection was only additional to the sound that was already present and much less critical.

3. CURVE FITTING APPROXIMATIONS.

3.1 Air Absorption.

Air absorption is a complicated function of distance, frequency, humidity and barometric pressure. For the purposes of this simple model, a single set of atmospheric conditions was assumed – 50% humidity at standard temperature and pressure.

From standard references, the attenuation for air absorption as a function of frequency was obtained. For a set of values of filter coefficient, the equivalent propagation distance which best fitted each filter response was found by a least-squares fit. Fig. 2 shows a typical result, for a relatively large distance of 68 m. For small distances, both the effects of air absorption and the errors due to the approximations are too small to provide reasonable illustrations of the effects.

A regression analysis was then used to find the best fit between the filter coefficients and the equivalent propagation distances. The resulting expression for filter coefficient, β, as a function of propagation distance, d, was found to be

$$\beta = 1.752 - 0.308 \log_e(d + 10.33)$$

Fig. 3 shows the 'fit' of this approximation function to the derived filter coefficients. Up to a distance of 50 m, the maximum error at any frequency was less than 0.3 dB. Beyond that, the maximum errors became progressively larger until, at 150 m, they were about 3 dB and, by 220 m, they were about 6 dB. However, those large distances would not be encountered very often and the overall attenuations would be large, making the subjective effect of such errors insignificant.

The maximum theoretical distance allowed by that approximation (β = 0) was 295 m. That was outside the range permitted by the room geometry, even for reflection path lengths (the maximum room size was 99.9 m in any direction). It was just possible to envisage a combination of largest permitted room with a large internal object giving a total path length of about 280 m. Those are not realistic conditions. In any case, the overall signal level would be nearly -50dB and essentially inaudible.

Overall, this was considered to be a reasonable approximation.

3.2 Reverberation Air Absorption.

The same model was used to calculate the reverberation filter coefficients as for free propagation. The only slight difficulty was to work out what the equivalent distance should be. Fig. 4 shows a simplified signal flow path for the reverberation processor [2]. In principle, the signal continues to circulate around a 'figure of 8' path indefinitely. In each complete loop, the signal passes through two filter stages. The decay rate (or reverberation time) was set by the filter gain and the 'air absorption' by the filter coefficient.

For the nominal scale reverberator[1], the delay round the whole loop was 451 ms at 48 kHz sampling frequency, corresponding to a physical distance of 153 m. Thus the equivalent distance for each of the air absorption filters was 76.67 m.

That same mechanism was also used to model three different types of reverberation. As expressed above, the damping would be appropriate for an empty room - one with uniform, wideband surface absorption to give the correct low-frequency reverberation time and only the air absorption to provide any frequency-dependency. That type of room was designated "Empty". In a broadcasting studio, the design target is usually a flat reverberation time - the acoustic treatment being tailored to have a falling absorption characteristic to compensate for the air absorption. That was implemented in the model simply by setting the damping to a very small value (0.01). That room type was designated "Studio". A third type of room was modelled with double the amount of normal air absorption to give a relatively steeply falling RT characteristic. It was designated "Furnished" to represent rooms with relative soft furnishings, especially carpets and drapes.

These three simple rooms types were included for demonstration purposes – there was clearly no reasonable limit to the number of different types of room characteristic which could be modelled,

[1] The reverberation process was limited to a relatively small range of useful reverberation times. To increase the range of available room sizes and reverberation times it was necessary to scale the reverberation process itself. Thus, both the delay around the loop and the air absorption filter coefficient were functions of the room size.

especially if the filter order were to be increased to allow more complex response shapes. One common type that would definitely need to be included in a final product is the domestic "living room", with its pronounced mid-band peak in the reverberation time.

3.3 Human Voice Directionality.

The modelling of an average human voice directionality was based on a very old reference [3]. The main purpose was to develop the dsp program rather than to model the responses exactly. Nevertheless, it is probable that the differences between this implementation and a more accurate one would not be significant in this application. No attempt was made to model the fine detail of the radiation pattern, which would in any case be different for different speakers.

As in all the other filter implementations, this one was also based on a single-stage, IIR lowpass filter. Data for the frequency responses of a human voice were extracted from Ref. 3. The printed graph was very small and only contained results for four frequencies. Data for those frequencies was extracted (by inspection) from the graph at 15° intervals. Null data for a very low frequency (20 Hz) was added to the data set to improve the curve fits, on the assumption that the curves should all pass through 0 dB at 0 Hz.

The filter coefficient and attenuation values were derived, for both α and β, by a multi-dimensional, curve-fitting procedure which minimized the sum of the squared errors for all 5 frequencies and 13 angles simultaneously. The resulting functions were: -

$$\beta = 1.225166 - 0.220645 \log_e (|\, angle|)$$
$$\alpha = -0.011433 \,|\, angle|$$

where $3° \leq |\, angle| \leq 180°$ and *angle* was measured relative to the 'straight ahead' orientation. Angles less than 3° would have produced filter coefficients greater than unity, if not otherwise allowed for. Such angles were considered to be effectively 0°. Fig. 5 shows, as an example, the source data and the approximation for an angle of 135°.

The resulting errors were relatively large, up to about ±2 dB over the whole range of frequency and angle. A significant proportion were randomly scattered and undoubtedly due to errors in reading of the small graph. At angles greater than ±90°, the errors were slightly less, ±1.0 dB. These functions will probably give values close enough to the correct ones to be subjectively identical. The errors will almost certainly make no perceptible difference to the overall result. The functions could easily be refined, if necessary, using more complete and accurate data.

It was also necessary to correct the reverberation level and frequency response for the source directivity index. The reverberation module input filter was used to set the directivity index frequency response. A least-mean-squares curve-fitting procedure for the directivity index from the reference produced a best fit value for the filter coefficient of 0.72 and a value for the overall attenuation of 1.2 dB. Those parameters were included in the setting of the reverberation module for 'human voice' directionality if appropriate. Of course, those values did not depend on any angles or distances.

3.4 Diffraction.

Fig. 6 shows a sketch of a direct sound path obstructed by an obstacle. It also indicates the two dimension parameters, offset *y* and normal source-object distance *l*, important for the calculation of diffraction. The obstacle is shown at right angles to the line from source to receiver. In practice, it might be at any angle. However, the theoretical effects of diffraction are not affected very much by

that angle and the calculations were carried out as though the obstacle was always perpendicular to the line from the edge of the obstacle to the receiver.

The calculation of the obstruction or reflection parameters required a number of steps. First, the coordinates of the source, the obstacle and the receiver were translated and rotated so that the source was at the origin of the coordinate system and the obstacle parallel to the *x*-axis. Fig. 7 shows the geometry in the transformed coordinate system. A flag indicating a potential obstruction or reflection condition was set, depending on whether the receiver and source were on the same or opposite sides of the obstacle line.

The processing then depended on whether there was a potential reflection (source and receiver on the same side of the obstacle) or not. If there was, then the source image position was calculated (trivially, on the *y*-axis, twice as far from the source as the obstacle). Then the intersection, I_r, was calculated between of the line of the obstacle and the line from source image to receiver. If that intersection fell within the span of the actual obstacle coordinates then a real (geometric) reflection was present. The source image coordinates and the intersection point, I_r, were then transformed back to real-world coordinates to allow the calculation of true distance and angle for the reflection. The system was slightly refined to provide information about near-reflections, so that the reflection source could be faded in progressively over a small range of near-geometrical reflections. That was necessary to avoid sudden discontinuities in the signal switching.

If there was potential obstruction then the intersection, I_o, was calculated between the line of the obstacle and the line from source to receiver. If that intersection fell within the span of the actual obstacle coordinates then a real (geometrical) obstruction was present. In that case, the real obstacle was extended somewhat (10% at each end) to allow for the non-geometric acoustic obstruction. This was done as a pre-screening step only - the final calculation was done later. The geometry was then further processed to obtain the parameters *y* and *l* for the calculation of diffraction. That same calculation procedure also had to decide finally on the full/partial obstruction conditions to set appropriate flag values.

An analytical study of diffraction at an edge involves half-order Bessel functions (as solutions to the wave equation in cylindrical coordinates) and is self-evidently too complex to incorporate in the acoustic model in that way. It is also complicated by a discontinuity at the boundary of the geometrical shadow zone.

The curve fitting for diffraction was by far the most complex of the approximation procedures. It is not feasible to describe it in a completely comprehensible way in this paper, so that what follows is even more of an outline than for the other approximations.

First, the theoretical solutions to the acoustic problem (Fresnel functions) had to be expressed in a numerical form [4]. That was done using standard textbook approximations to the functions. The results depended on *y/l* (see Fig. 6) rather than *y*, so *y/l* became the independent parameter for the offset. Fig. 8 shows, as an example, the sound pressure level as a function of the scaled offset, *y/l*, for a frequency of 1080 Hz, calculated using the numerical approximations to the Fresnel functions[2]. The amplitude ripples of the Fresnel zone are clearly visible and, equally clearly, difficult to find approximations to. At that stage, the amplitude responses were functions of geometry and not of

[2] The apparent irregularities of all of the responses shown in figures 8 – 13 are mainly 'features' of the spreadsheet used for the calculations, compounded in some cases by the coarse frequency resolution (1/3rd octave bands). The figures are all for illustration only – all calculations were carried out to adequate, multi-digit accuracy.

frequency. To match filter characteristics it was necessary to express the responses as functions of frequency for fixed geometrical arrangements.

The numerical expressions were re-written to include frequency as an independent parameter and the results were graphs like Fig. 9 and Fig. 10, which show responses as functions of frequency for fixed values of y/l and l. Fig. 9 illustrates a typical response in the Fresnel zone ($y/l = +0.3$) and shows the expected ripples. Even so, the ripples can be seen to be relatively shallow. Fig. 10 is for the full shadow zone ($y/l = -0.3$) and is monotonic, showing none of the cyclical variations characteristic of the fringing zone.

Very many such graphs were produced (210 in all) for large ranges of the two parameters l and y/l. The results formed a 4-dimensional data set for 30 values of y/l and seven values of l. For each of those combinations, the result was a frequency response function which was to be modelled by a combination of attenuation and single-stage, IIR lowpass filter.

For each frequency response, a least-mean-squares regression was carried out to find the best-fit pair of filter coefficients, α and β. The errors involved in that process were around ±2 dB for positive and small negative values of y/l (-0.5 to 0.5) rising to ±4 dB for negative values of y/l like -2.0. Whilst these were relatively large values of error, they corresponded to errors of frequency response rather than of overall level and the larger errors occurred when the signal itself had been attenuated very significantly (more than 20 dB). It is not likely that these frequency response errors would be perceptible. In any case, with simple filter topologies and curve fitting, little else is possible.

That curve-fitting process produced a pair of numerical values for each of the 210 combinations of y/l and l. Some form of approximation then had to be found to generate those values from the basic geometrical parameters. Because of the discontinuity at $y/l = 0$, the curve fitting had to be carried out in three sections.

The final results were: -

for most of the fringing zone ($y/l \geq 0.05$)

$$\beta = 1$$

$$\alpha = a \log_{10}(y/l) + b$$

where: -

$a = -0.091475\, d^2 + 0.238897\, d + 3.543441$

$b = -0.062588\, d^2 + 0.514117\, d + 0.653043.$

$d = \log(l) / \log(2)$

The maximum errors in the fitting of α to the response approximations over the range $0.5 \leq l \leq 32$ and $0.05 \leq y/l \leq 0.5$ were +0.22 dB and -0.24 dB.

for the near fringing zone ($0.05 \geq y/l \geq 0.00$)

It was impractical to find a curve-fit for positive values of y/l between 0.0 and 0.05. For that range, a simpler linear interpolation between values for the fringing zone at $y/l = 0.05$ and $y/l = 0$ was used: -

$$\alpha_0 = 0.0432\, d^2 - 0.441\, d - 4.6261$$

$$\alpha_{0.05} = 0.0719\, d^2 + 0.096\, d - 3.908$$

For values of $y/l > 0.5$ the effects of diffraction are very small. In fact, the diffraction limit was set so that such values were not treated as obstruction at all.

for the shadow zone ($y/l \leq 0$)

$$\beta = a / (b - y/l)^c - 0.066$$

where: -

$a = -0.034725\ d + 0.305212$

$b = 0.013243\ d + 0.875066$

$c = 0.103434\ d^3 + 0.06790\ d^2 + 0.2.065981\ d + 6.381259$

and with d as above.

The maximum errors in the fitting of β to the approximations over the range $0.5 \le l \le 32$ and $-3.0 \le y/l \le 0.0$ were +0.012 and -0.012 (very small indeed – certainly not subjectively perceptible in such an acoustic model, though they might be in careful *A/B* tests of a filter alone).

$$\alpha = a + b / (c - y/l)^{0.6}$$

where: -

$a = -3.745552\ d - 17.63651$

$b = -0.239024\ d^2 + 2.285254\ d + 11.78731$

$c = -0.092269\ d + 0.84511$

The maximum errors in the fitting of α to the response approximations over the range $0.5 \le l \le 32$ and $-3.0 \le y/l \le 0.0$ were ±0.5 dB. As an example, Fig. 11 shows the best-fit approximations for α in the shadow zone.

For $y/l < -3.0$, β tended towards 0.066 as did the approximations. For $l > 32$ m, the errors remained small (< 1 dB). Both of those functions could therefore be extrapolated beyond those limits, if necessary. However, such extreme values would be rarely encountered in practice.

Figs. 12 and 13 show, as examples, the final filter responses for two cases, one in the Fresnel zone, $l = 1$, $y/l = .1$, and one in the shadow zone, $l = 1$, $y/l = -2$. (The solid lines represent the actual filter response and the dashed lines represent the intended response, obtained from the numerical approximations to the Fresnel functions). In the Fresnel zone (Fig. 12), the desired response is fairly uniform, with mostly overall attenuation. The slight frequency irregularities were not reproduced at all – the filter simply consisted of wideband attenuation. The maximum errors were around 2 dB. In the shadow zone (Fig. 13), the errors amount to ±6 dB, with insufficient low frequency and high frequency responses and an excessive midband response. However, the entire signal is between 10 and 40 dB relative to the direct path, which would itself be further attenuated because of the total distance from source to receiver around the sides of the object. It is highly improbable that even such relatively large errors would result in the illusion being perceived as 'incorrect', especially when accompanied by the visual cues.

4. PERFORMANCE.

4.1 Overall Quality.

The static noise and distortion (THDN) performance of the complete model were around -78 dB relative to the reference level (22 Hz - 22 kHz unweighted). The dynamic range of the system extended to about +18dB above that, but was somewhat dependent on the model geometry and the source signal[3]. Most of the noise and distortion arose from the fixed-point, 24-bit quantisation limit

[3] The internal signals were attenuated by 6 dB to allow for some additional headroom in the dsp system. The signals were processed throughout the system without any additional attenuation, until the final step of combining them into the five loudspeaker drive signals. Because of the addition of many intermediate signals, especially in the reverberation processor, there was a significant potential for very large signal levels to occur

but was judged to be adequate (at least for an experimental system). The actual distortion and noise floor of the dsp system itself, measured using a simple mixer example supplied by the manufacturer, was not measurably different (±0.5dB). That showed that the large amount of signal processing (at 24-bit resolution) did not significantly affect the 18-bit output.

Tests carried out on individual component signals showed that all parameters appeared to be implemented according to the intended functions. By the nature of the system, it was not practicable to measure the responses in detail for every possible combination of system parameters.

4.2 Subjective Impression.

Subjectively, the model appeared to be too reverberant. This aspect was discussed in the original paper [1] and is a feature of all virtual audio systems. It is well known (and common experience) that a space does not give the same subjective impression when the listener is actually in the space, compared with a recording of the same space. This psycho-acoustic factor probably applies to all of the subjective impressions – not just to the reverberation.

The subjective impressions given by some of the other model features were also surprising, though as far as possible it had been established that the system was working as intended. For example, the effects of an obstruction were not as dramatic as had been expected. Although there was a pronounced loss of the direct signal and change in the frequency response as the source was obstructed it was not as much as had been expected. In reasonable sizes of rooms, the overall signal quality remained largely governed by the early reflections and reverberation, to a surprising extent (at least to this author). It is quite likely that the same sort of psycho-acoustic effects as occur for reverberation apply to other sorts of effects as well. These differences between being present in a real room and listening to a model or recording have implications for many kinds of acoustic modelling.

It is probable that many aspects of the acoustic model will have to be modified from their objectively correct values, either moderated or exaggerated, to produce the 'correct' subjective effect. This can only be determined by tests *with accompanying pictures*. That will require experimental programme productions, which, at the time of writing, have not yet been carried out.

5. SUMMARY AND CONCLUSIONS.

This paper has attempted to summarise the implementation of simple filter functions for an audio processor for creating virtual sound fields.

All of the frequency response characteristics were achieved by approximations using simple, first-order IIR filters with only two control parameters – overall gain and feedback fraction. The approximations were calculated by least-squares regression curve-fitting using standard, reference propagation functions and the theoretical filter responses. Expressions were then derived for the relationships between the filter control parameters and the acoustic parameters. For example for air absorption, the two filter parameters were expressed in terms of propagation distance for fixed atmospheric conditions.

In most cases, the approximations provided reasonably close or even good fits to the physical effects being modelled. In some others, the response errors were substantial, especially for extreme

---

at some frequencies. In practice, because of the statistical incoherence of the large number of contributions, such overloads did not occur and 6 dB was found to be a reliable working margin, even for sinusoidal signals.

conditions. Even in those cases, however, the responses obtained for most useful conditions were good enough to provide a convincing audio illusion, especially in conjunction with the accompanying pictures. The largest errors occurred for conditions where the sound would be greatly attenuated anyway - for example, air absorption at distances greater than 100 m.

Greater objective realism could undoubtedly be achieved with more complex filter topographies, at the cost of reductions in processing speed, but it is doubtful whether additional objective realism would significantly enhance the overall subjective impression.

Although the room model was otherwise empty, one internal object could be included, if required, in order to model obstruction and reflection by large sections of the studio set. That had, perforce, to include a model of the diffraction of sound around an obstacle. The modelling of the diffraction responses was particularly complicated.

The system appeared to perform adequately and objective tests showed that it was behaving as calculated and that the measured noise and distortion performances were essentially indistinguishable from those of the basic dsp system.

Subjective tests, using experimental productions, will have to be carried out to assess the operating parameters and to optimise the control functions. It is probable that this development will take place as the system is used for some types of productions. It has also become clear that the system may have applications in fields other than virtual production, such as post-production, film dubbing and even radio drama. Whilst it is true that all of the effects can be realised by existing equipment, it is generally too expensive (mainly in operator time) to attempt to create reasonably realistic artificial acoustic environments by such *ad hoc* means. It is also difficult to track movement in that way. It may well prove useful to use such a generator of synthetic acoustic environments to assist in such productions.

6. REFERENCES.

1. Walker, R. A simple acoustic room model for Virtual Production. Reproduced Sound 15 Conference, Stratford, November 1999.
2. Dattorro, J., Effect design. JAES, **45**, 9, 1997, September.
3. Dunn, H.K. and Farnsworth, D.W., Exploration of pressure field around the human head during speech. J.Acoust. Soc. Am., 10, 184-199, 1939.
4. Morse and Ingard, Theoretical Acoustics. Princeton University Press, ISBN 0-691-02401-4, p.450.

7. ACKNOWLEGEMENTS

This paper is published by permission of the British Broadcasting Corporation.

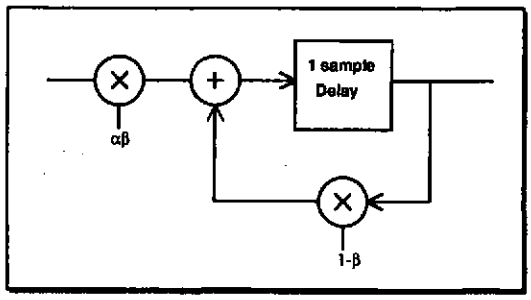

Fig. 1. Basic filter arrangement.

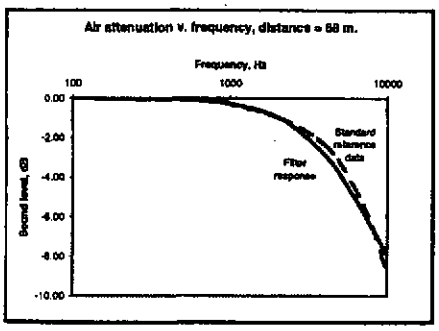

Fig. 2. Relative air absorption response for a distance of 68 m.

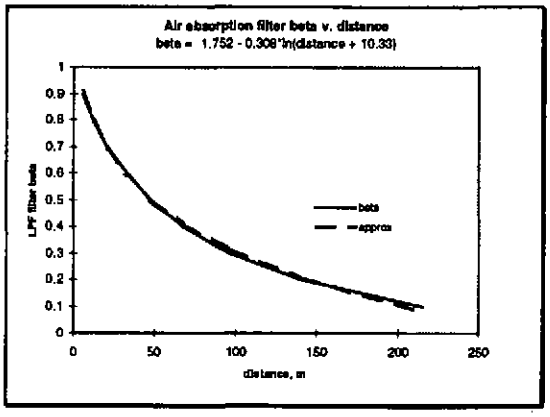

Fig. 3. Air absorption filter beta as a function of distance.

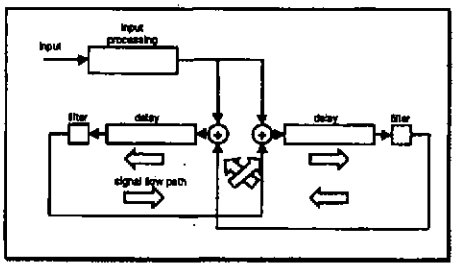

Fig. 4. Basic reverberator signal paths.

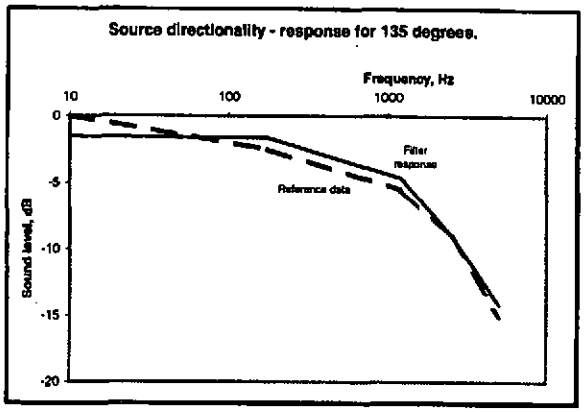

Fig. 5. Source directionality, response at 135 degrees.

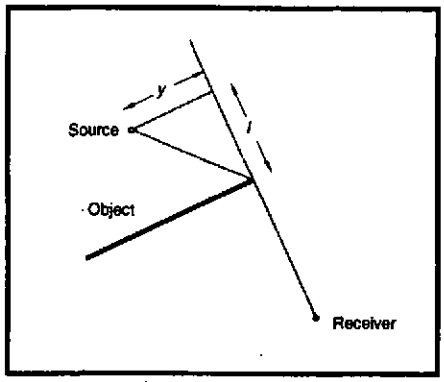

Fig. 6. Geometry of a fully-obstructed source.

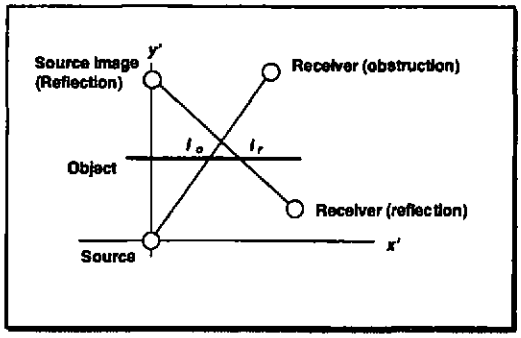

Fig. 7. Transformed coordinates for obstruction/reflection calculation.

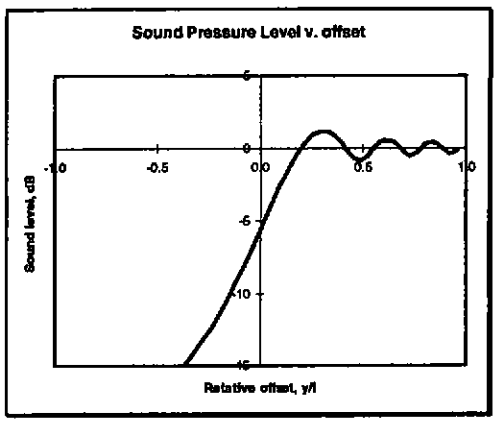

Fig. 8. Example of sound pressure level as a function of the scaled offset, *y/l*, for a frequency of 1080 Hz.

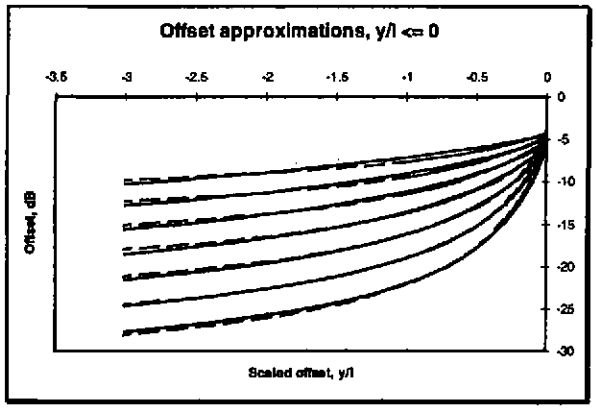

Fig. 11. Diffraction offset approximations for shadow zone.

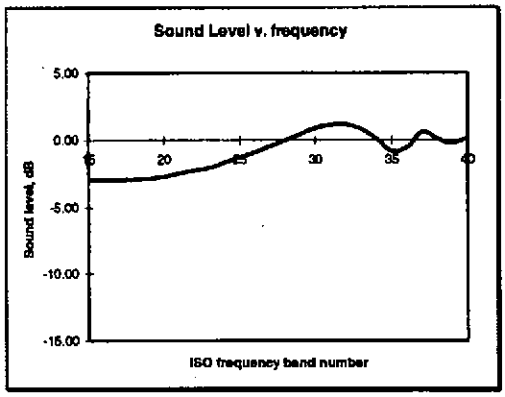

Fig. 9 Theoretical diffraction amplitude response for fringing zone (*l* = 2, *y/l* = 0.3).

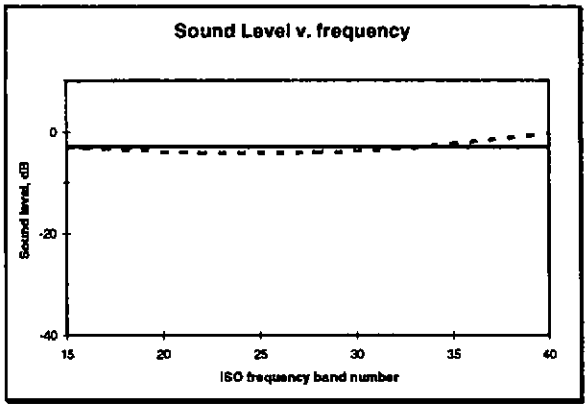

Fig. 12. Approximated diffraction response in Fresnel zone, l = 1, y/l = 0.1.

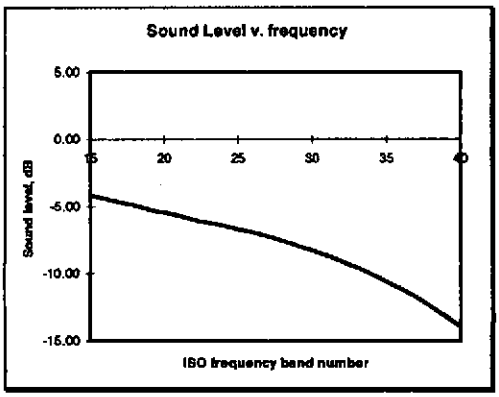

Fig. 10. Theoretical diffraction amplitude response for shadow zone (*l* = 2, *y/l* = -0.1).

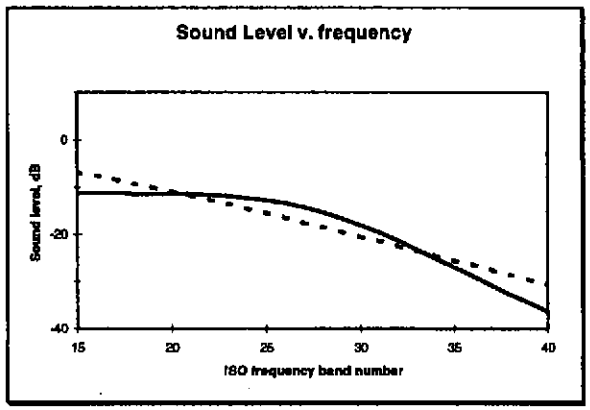

Fig. 13. Approximated diffraction response in (deep) shadow zone, l = 1, y/l = -2.