Audio Engineering Society
Convention Paper
Presented at the 126th Convention
2009 May 7–10 Munich, Germany

Temporal matching of 2D and 3D wave-based acoustic modeling for efficient and realistic simulation of rooms

Jeremy J. Wells[1], Damian T. Murphy[2], and Mark J. Beeson[3]

[1,2,3] Audio Lab, Department of Electronics, University of York, York, North Yorkshire, YO10 5DD, England
jjw100@ohm.york.ac.uk, dtm3@ohm.york.ac.uk

ABSTRACT

Methods for adapting the output of a two dimensional Kirchoff-variable digital waveguide mesh to better match that of a 3D mesh, both of which are intended to model the same acoustic space, are presented. Details of the methods, including quality of output and computational demands, are given along with details of how they are incorporated into the hybrid system within which they are employed.

1. INTRODUCTION

The simulation of the acoustics of real and imagined environments using digital techniques has been an active area of audio research for almost half a century. Earlier digital reverberators typically employed networks of one dimensional finite impulse response (FIR) and infinite impulse response (IIR) filters [1]. However, in recent years approaches which explicitly consider the geometry and materials of the space being simulated have received greater attention as they become more computationally tractable and the numerical methods required to implement them have been improved and better understood.

There are two main approaches to modeling an acoustic space in two or three dimensions: ray-tracing and wave-based methods. The former models the propagation of sound in a space as straight line 'rays' of sound that emanate from the sounding object and are reflected, diffused and attenuated by objects with which they collide. This geometric approach is relatively cheap computationally and is effective at mid and high frequencies where the size of objects that the sound rays interact with are large compared to the wavelength of the sound that the rays represent. However, at low frequencies sound wave propagation is not well modeled by rays. The second approach models the sound directly as variations in pressure and/or velocity which propagate in two or three dimensions. Although

this approach is computationally intensive, actual wave propagation effects such as diffraction are inherent in the model. These effects are important at low frequencies where the size of objects within in the room, and in many cases the room itself, are comparable to the wavelength of the sound being modeled. Since rays are straight lines of infinitesimal width, effects such as diffraction are not inherent since the rays are not able to 'bend around' objects.

RenderAIR is a room acoustics modeling tool which uses a combination of geometric (ray-based) and wave-based simulations of sound propagation in enclosed spaces [2]. This affords the user very fine control over the trade-off between model accuracy and computational efficiency while combining the most desirable aspects of both wave and ray based modeling. It is a cross-platform system with applications in sound design for music and film as well as architectural design. It offers an integrated development environment within which acoustic spaces can be designed from scratch or imported using common interchange formats for 3D applications (such as Collaborative Design Activity – COLLADA). Once a desired virtual space has been designed or imported and source and receiver(s) have been positioned, an impulse response is rendered off-line which can then be used via a real-time or off-line convolution algorithm. Work parallel to this has focussed on developing sources and receivers with first- or second-order directional characteristics from their omnidirectional counterparts [3], [4].

The hybrid approach of *RenderAIR* requires the combination of outputs from three different types of model: two dimensional (2D) and three dimensional (3D) wave-based and 3D ray-traced. Typically a band and time limited 3D output is used to provide the best possible model of the low frequency, early time behaviour of a space in response to an impulse. In both the natural and computational domains, 2D and 3D reverberators do not behave in exactly the same way, particularly with regard to reverberation time and, to a lesser extent, magnitude response. Before the outputs from 2D and 3D models are combined the 2D impulse response is modified in order to ensure that it better matches the 3D model and the response of the actual space being simulated.

The rest of this paper presents details of the algorithm used to modify the 2D response before combination with the 3D wave-based and ray-traced outputs. In the next section an overview of the digital waveguide mesh (DWM) and finite-difference time domain (FDTD) methods for the modeling of wave propagation is given. In Section 3 the underlying acoustic theory pertaining to this work is summarised. Section 4 gives details of the 2D-3D correction algorithm. Section 5 provides a comparison of the outputs from 3D, 2D and '3D-adapted 2D' meshes for spaces of varying size and shape, diffusive and absorptive properties.

2. WAVE-BASED MODELING WITH THE DIGITAL WAVEGUIDE MESH

For wave-based simulation in both two dimensions (2D) and three dimensions (3D) *RenderAIR* uses the digital waveguide mesh (DWM). This is a multi-dimensional arrangement of nodes and unit delays. In what is known as a K-variable (K-) DWM, the nodes (which sum and weight pressures arriving from adjacent nodes) and unit delays (which connect the nodes and represent the propagation time/distance between them) are arranged to form a mesh which represents the physical space being modeled [5]. The K-DWM is an example of a finite difference time domain technique (FDTD) for solving the 2D and 3D wave equation. The distance (in the virtual space) between the nodes is determined by the sampling rate of the mesh – a high sampling rate requires a dense mesh and so a large virtual space requires a large number of nodes and hence a large number of nodal updates per sample. Also, in practice there is a dispersion error associated with the DWM due to a direction dependent variation in wave propagation speed due to the topology of the mesh [6]. This error increases with frequency and requires either over-sampling of the mesh or frequency warping of the output. Therefore, for an accurate full bandwidth rendering of the impulse response of a particular space in three dimensions there is a computational expense (both in terms of memory to store the mesh and the speed to update it) which is prohibitive in most situations. It is for this reason that *RenderAIR* offers the possibility of a hybrid impulse response, allowing the user to trade computation time with model accuracy.

2.1.1. 2D versus 3D wave-based modeling

Natural reverberation typically occurs within 3D resonators although, historically, some use has been made of plates which are essentially 2D resonators. Modeling of acoustic spaces with a 2D model is many times faster in computational terms than 3D modeling.

For example, considering a 1 meter cubic room in 3D the number of nodes required at a sample rate of 44.1 kHz is given by:

$$N = W \times L \times H \tag{1}$$

where *W, L* and *H* are the numbers of nodes in the width, length and height direction. The number of nodes in a given direction is related to the inter-nodal spacing which is determined by the type of mesh being used. For example, the number of nodes needed to span the width of a mesh of the type used in *RenderAIR* is given by:

$$W = \text{round}\left(\frac{F_S w}{c\sqrt{D}}\right) + 1 \tag{2}$$

where w is the width of the room in meters, c is the speed of sound in air (assumed to be 344 m/s in this paper), D is the number of dimensions (three in this case) and F_S is the sample rate (sometimes referred to as the 'update frequency') [5]. In this example the number of nodes required in 3D is 421,875. For a 2D model the number of nodes is given by:

$$N = W \times L \tag{3}$$

where the number of nodes in each dimension is calculated using (2) with D=2. For this example the number of nodes required in 2D is 8,464 if the 2D plane is parallel to the floor/ceiling (or either pair of facing walls) or 16,641 if, at the other extreme, the 2D plane is at an angle of 45 degrees to one of the surfaces. When translating from a 2D to a 3D model there is an inherent ambiguity in the way in which the 2D plane intersects the 3D space that it is approximating. In *RenderAIR* this is resolved by having the plane intersect both the source and receiver position. There is still rotational ambiguity but this is usually resolved by having the plane flat (i.e. so that if the source were a human listener it would intersect both ears).

The relationship between F_S and the effective upper frequency of the mesh is determined by the type of mesh employed. *RenderAIR* uses a K-mesh rectilinear (also known as 'standard leap-frog') scheme where, at each time-step, an air node (i.e. one that is not next to a boundary on any of its sides) is updated according to the following equation for 2D:

$$\begin{aligned} p_{l,w}^{n+1} &= \lambda^2 \left(p_{l-1,w}^n + p_{l+1,w}^n + p_{l,w-1}^n + p_{l,w+1}^n \right) \\ &+2\left(1-2\lambda^2\right) p_{l,w}^n - p_{l,w}^{n-1} \end{aligned} \tag{4}$$

and the following for 3D:

$$\begin{aligned} p_{l,h,w}^{n+1} &= \lambda^2 \begin{pmatrix} p_{l-1,h,w}^n + p_{l+1,h,w}^n + p_{l,h-1,w}^n \\ +p_{l,h+1,w}^n + p_{l,h,w-1}^n + p_{l,h,w+1}^n \end{pmatrix} \\ &+2\left(1-3\lambda^2\right) p_{l,h,w}^n - p_{l,h,w}^{n-1} \end{aligned} \tag{5}$$

where p is pressure, l, h and w are now integer indices in each spatial direction and n is the current sample number. In both equations λ is the Courant number that relates the distance between samples in space and time and the speed of wave propagation:

$$\lambda = \frac{cT}{X} \tag{6}$$

where T is the time sampling interval and X is the spatial sampling interval [6]. By choosing $\lambda = \sqrt{1/2}$ for 2D and $\lambda = \sqrt{1/3}$ for 3D equations (5) and (6) simplify and meshes can be updated using an alternating sub-gridded scheme where only half of the nodes need to be updated at each time-step provided the signal with which the mesh is excited is itself sub-gridded (i.e. every other sample is zero) a condition which is a met for a unit impulse [6]. Such a scheme offers significant saving both in terms of computational cost and the memory required to accommodate the mesh. The disadvantage of such sub-gridded schemes is that their effective upper frequencies are often lower than other types of mesh update scheme, so there is a tradeoff between bandwidth and computational cost to be considered. Comparisons with other schemes have indicated that the sub-gridded rectilinear scheme adopted by *RenderAIR* performs well in terms of model accuracy and efficiency, when over-sampled to compensate for the reduced effective upper frequency.

3. REVERBERATION TIME IN 2D AND 3D

The prediction of the reverberation time of a room from its geometry, contents and the materials from which it is constructed is a problem that was famously first tackled in earnest by Wallace Sabine, over a century ago [7]. Sabine derived a formula relating room volume (V), surface area (S) and average absorption coefficient (α)

to reverberation time from a series of measurements which showed good agreement with that observed for rooms with low overall absorption. Subsequent work in this area focussed on finding an analytic, rather than empirical, solution to reverberation prediction. This work attempted to determine the mean free path that a sound wave would travel between reflections and the amount of absorption that would occur at each reflection. From this the time taken for the sound energy associated with a single wave, or all of the waves propagating in a room, to fall to a certain level could be predicted. Such an approach is the basis of the Norris-Eyring formula, of which the Sabine formula of thirty years earlier is a very good approximation where the average absorption of the room is not greater than α = 0.5:

$$RT_{60,3D} = \frac{-4V \ln\left(10^{-6}\right)}{cS \ln(1-\bar{\alpha})} \quad (7)$$

where $RT_{60,3D}$ is the time taken for the sound energy to fall to one millionth of its initial value (-60 dB) and $\bar{\alpha}$ is the absorption coefficient of the room averaged over all surfaces and all frequencies [8]. The mean free path is given by $4V/S$ and, in some derivations of this quantity (e.g. that presented in [9]), is independent of room shape. However it is important to note that it is assumed that the room is perfectly diffuse, with no specular reflections from surfaces, and the absorption of all surfaces is the same (i.e. $=\bar{\alpha}$). Applying the same approach to the 2D case, the reverberation time of a plane is given by:

$$RT_{60,2D} = \frac{-\pi S \ln\left(10^{-6}\right)}{cC \ln(1-\bar{\alpha})} \quad (8)$$

where the upper-case C is the total length of all edges (i.e. the circumference of the plane and any objects within it) [9]. Absorbers typically exhibit frequency-dependent characteristics and the formulae given (7) and (8) can incorporate this by becoming measures at a specific frequency of the reverberation time given the average absorption at the same frequency. A generalization of these formulae to rooms with non-uniform absorption is given by the Millington-Sette equation:

$$RT_{60,3D}(f) = \frac{-4V \ln\left(10^{-6}\right)}{c\sum_{i}^{I} S_i \ln\left(1-\alpha_i(f)\right)} \quad (9)$$

$$RT_{60,2D} = \frac{-\pi S \ln\left(10^{-6}\right)}{c\sum_{i}^{I} C_i \ln\left(1-\alpha_i(f)\right)} \quad (10)$$

where I is the number of different absorber types and i is the index of the absorber with regard to either its surface area or circumference and its frequency dependent absorption.

4. 2D AND 3D REVERBERATION MATCHING

As already stated, it is often the case that a full bandwidth output from a 3D DWM is not computationally tractable either in terms of the memory physically available or the time which can reasonably allocated to the task (obtaining the output from a large 3D mesh with a high sampling frequency can take days). To this end the 3D DWM is combined with the output of a 3D ray-tracing algorithm and a 2D DWM. Typically the ray-trace and 3D DWM are combined to produce the initial (early reflections) part of the response, and the ray-trace and the 2D DWM are combined to produce the rest of the response (the reverberation).

Whilst the 3D DWM and the ray-trace are both based on the same virtual space, the 2D DWM is a virtual plane which is an approximation to that space and its output impulse response, specifically its RT_{60}, is likely to differ. This section describes a simple correction algorithm which is applied to the output of the 2D mesh to correct the RT_{60} so that it is the same as that of the 3D mesh.

A fundamental design goal in the development of *RenderAIR* is the fidelity of its output to that which is produced by the actual space that is being modeled (or, if the space does not currently physically exist, what would what be produced if it did). A useful by-product of this is that the virtual spaces it produces match their real counterparts in terms of their behaviour well, meaning that analytical methods applied to real 2D planes and 3D spaces, as outlined in the previous section, can be applied to these virtual spaces.

The purpose of the 2D/3D matching algorithm described here is to adjust the RT_{60} of the 2D output so that it matches that of the 3D and ray-traced output. Since the 2D output is only used to synthesize the latter part of the reverberation temporal alignment of the sparse, but temporally distinct, early reflections is not required. Beyond the 'early time transition' the reflections should be sufficiently dense in time that adjustment of the overall decay time is all that is required for temporal matching. Since the outputs of the all parts of the system are due to linear time-invariant processes this matching algorithm can be performed post DWM rendering which is much cheaper than attempting to incorporate it into the 2D model.

Since *RenderAIR* allows the absorption and diffusion characteristics of all materials used in the virtual room to be defined at octave intervals from 125 Hz to 8 kHz the correction algorithm also divides the spectrum of 2D output so that the RT_{60} can be corrected in each of these frequency bands. This is achieved by the use of a lengthily tapped linear phase orthogonal octave filter bank (353 taps for F_s = 44.1 kHz) where there is a low and a high pass filter in addition to band-pass filters centered at 125, 250,...8 kHz. The target magnitude response of the filters is zero gain at the octave above and below the centre frequency (although, of course, this cannot be achieved practically with a finite number of taps). The correction of the RT_{60} in each filtered band is achieved by applying a time-varying gain/attenuation to the output of each filter, before the outputs are summed to produce the full bandwidth 3D-corrected 2D output.

The gain is applied according to a linear function determined by the difference between the attenuation per sample exhibited in both the 2D and 3D impulse responses:

$$\Delta G = \frac{60}{F_S}\left(\frac{1}{RT_{60,2D}} - \frac{1}{RT_{60,3D}}\right) \qquad (11)$$

where ΔG is the change, in dB per sample, applied to the filter gain. In *RenderAIR* this time varying filter gain for RT_{60} correction is combined with that applied to simulate air absorption, which is also a frequency dependent effect. Where ΔG is positive care must be taken to ensure that this does not cause an apparent instability in the system (i.e. net growth in the output).

5. RESULTS FOR SIMPLE MODELED SPACES

To demonstrate how the matching algorithm works in practice, results for some simple, but contrasting virtual spaces, are presented. For each of these spaces the 2D and 3D output along with the 3D-compensated 2D output DWM are compared. For ease of comparison both types of output have been rendered at F_s = 44.1 kHz. For the 3D scheme used here this gives an effective upper frequency of 16 kHz, for 2D it is 22.05 kHz. The meshes are excited with a unit impulse and the output is then processed to correct for the frequency response of the mesh, as outlined in [10], and to remove the reflected spectrum above the effective upper frequency. This is equivalent to exciting the mesh with an identically conditioned input signal. Where the surfaces of the spaces/edges of the planes have diffusive properties these are uniformly distributed throughout. The implementation of boundary conditions is based on work presented in [11]. For each 3D room the source and receiver are placed in diametrically opposite corners, which gives maximum spacing between them and ensures that all room modes are captured. For the equivalent 2D model the source and receiver are on the same plane which is parallel with the floor.

5.1. Cubic room, 4.7 x 4.7 x 4.7 m, non-frequency-dependent absorption

Rooms having all three dimensions the same length are the simplest geometrically, but are amongst the most problematic acoustically since they are modally degenerate. For this comparison the normal-incidence reflectance (R_0) of all of the surfaces is 0.9. The absorption coefficient is calculated from the reflectance at 45 degrees:

$$\alpha = 1 - R_{45}^2 \qquad (12)$$

where:

$$R_{45} = \frac{\xi_w \cos^{D-1}(45) - 1}{\xi_w \cos^{D-1}(45) + 1} \qquad (13)$$

and D is the number of dimensions (for example, in 3D there is an incident azimuth and elevation angle) and the specific wall impedance is given by:

$$\xi_w = \frac{1+R_0}{1-R_0} \quad (14)$$

This gives an estimated α of 0.26 in 2D and 0.34 in 3D.

The dimensions of the 2D plane are 4.7m x 6.7m. Using (7) the $RT_{60,3D}$ is estimated as 0.30s and, using (8) the $RT_{60,2D}$ is 0.58. This gives ΔG = -0.0022 dB/sample. Figure 1 shows the normalised energy decay for the 3D, 2D and corrected 2D impulse responses.

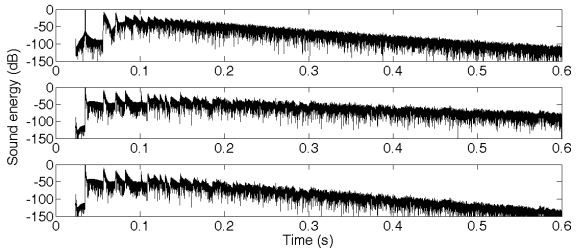

Figure 1: Energy plots of 3D (top), 2D (middle) and corrected 2D impulse responses for cubic room.

It can be seen that the slope of the energy decay for the corrected 2D response matches that of the 3D decay much better than the uncorrected 2D. The RT_{60} of each of these responses, extrapolated from measurements of the decay between -5 and -35 dB (ref. the initial energy), are: 0.34s (3D), 0.56 (2D) and 0.3 (corrected 2D). The corrected $RT_{60,2D}$ is a significant improvement, but the match is not perfect as (7) has under-estimated the $RT_{60,3D}$ and (8) has over-estimated $RT_{60,2D}$. One possible way of ameliorating the effect of mis-estimation of the $RT_{60,2D}$ would be to measure this from the actual 2D response, however this would break the equivalence between the two measures used to scale this response, and would need to be tested across a wide range of spaces in order to ensure that it actually improves the match between the 2D and 3D outputs in general.

5.2. Room with prime dimensions

A room with similar volume to the cubic room but with prime dimensions of 211 x 331 x 449 inter-node spacings, which corresponds to 2.9m x 4.5m x 6.1m, is tested next. Rooms with prime dimensions are desirable in acoustic terms since they will not have coincident modes (where a mode due to one path has the same centre frequency as a mode, or modes, due to another path). Again, R_0 = 0.9 for all surfaces giving α = 0.26 in 2D and 0.34 in 3D.

From (7) the $RT_{60,3D}$ is predicted to be 0.26s. For a 3D space with non-equal dimensions there is ambiguity in how to position the 2D plane as described in section 2.1.1. Whilst this ambiguity can be resolved by deciding which surface represents the floor of the space, it is instructive here to consider all three cases. This gives three 2D planes, each of which might be used to represent the 3D space. Their dimensions are 6.7m x 4.5m, 7.5m x 2.8m (the discrepancy between the 2.9m dimension for the 3D case is due to the fact that distances must be rounded to correspond to the nearest whole number of nodal spacings) and 6.1m x 5.4m. The $RT_{60,2D}$ for each of these planes are 0.57s, 0.44s and 0.60s respectively. From these estimates ΔG = -0.0029, -0.0022 and -0.0030. The measured $RT_{60,2D}$ are 0.54s, 0.42s and 0.57s. The corrected $RT_{60,2D}$ are 0.29, 0.26 and 0.26 showing good correspondence with the measured $RT_{60,3D}$ and, although there is some deviation in the first case, in all cases the temporal match between all impulse responses has been improved.

5.3. Cubic room, 4.9 x 4.9 x 4.9m with highly diffusive surfaces

There are two approaches to producing diffusive boundaries in DWMs. The first uses a W-mesh diffusing layer which interfaces with an interior K-mesh. A W-mesh has bi-directional delay lines allowing impedance mismatches at nodes to be explicitly modeled in the update equation [5]. The diffusing layer achieves diffusion by measuring the angle of incidence of the sound energy approaching a boundary (since a bi-directional delay line can have a direction associated with the propagation of pressure) and controlling the specularity of the reflection by adding a random offset to the angle of reflection. For a perfectly specular reflection no noise is added, for a perfectly diffuse reflection the angle of reflection is chosen entirely randomly, with no weighting given to the angle of incidence in the reflection decision. Control of the relative weighting of the angle of incidence and the random term allows control of the diffusive properties of the surface. The alternative approach models the physical cause of diffusion which is the variation in the position of boundary surface, rather than the effect itself by incorporating a delay into the wall impedance filter of the boundary [11],[12]. This is the approach adopted here since it does not require a W-mesh layer between

the inner and boundary nodes. The reflectance, and therefore absorption, are the same as for previous examples.

The measured and predicted $RT_{60,3D}$ are 0.40s and 0.31s and it is interesting to note that there is less agreement between these two values than in the non-diffuse case, whereas the theory of reverberation time prediction suggests that the two should be closer where the actual space is more diffuse. The situation is the same in 2D where the measured and predicted $RT_{60,2D}$ are 0.69s and 0.59s respectively. However, both the 2D and 3D estimation are affected in the same way; they are both under-estimated to a greater extent than in the non-diffuse case. The measured $RT_{60,2D}$ of the corrected response is 0.30, an improvement on the un-corrected case, but not as good as that for the non-diffuse cubic room.

5.4. Room with frequency dependent porous absorption

Acoustically porous materials absorb sound by damping the movement of air particles. This damping only occurs where the particles have velocity and so they have the greatest effect when they are placed a quarter-wavelength from a boundary. A method of predicting the frequency dependent reflectance of a porous absorber from its porous layer thickness and flow resistivity is given in [13] and adopted here to determine the octave band reflectance. These octave band reflectance points are then used to derive a 23rd order recursive-filter with the best least squares fit to the frequency response, from which the wall-impedance for the frequency dependent boundaries of the boundary is derived, as described in [11]. The frequency-dependent reflectance, and that of the digital filter used to model it, is shown in Figure 2, and can be seen that there is a very close match between the two.

Figure 3 shows the actual $RT_{60,3D}$ and $RT_{60,2D}$ in each of the octave bands. As can be seen there is a significant discrepancy between the values for the 3D case and the three possible sets of values for the 2D case (values for each of the possible planes are presented). Figure 4 compares the same $RT_{60,3D}$ with the $RT_{60,2D}$ of the time-corrected 2D responses. It can be seen that, although there are still some significant discrepancies, for individual responses at certain frequency points, the temporal alignment process has produced a better match overall between $RT_{60,3D}$ and $RT_{60,2D}$.

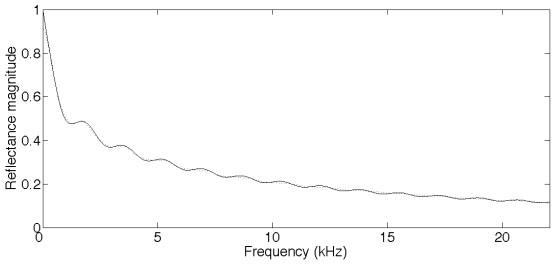

Figure 2: Magnitude of frequency-dependent reflectance for a porous absorber (solid line) and magnitude response of equivalent 23rd order recursive digital filter (dotted line).

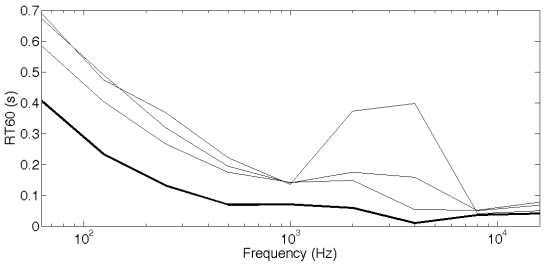

Figure 3: Measured $RT_{60,3D}$ (thick line) and $RT_{60,2D}$ (thin lines) for a DWM with porous absorber.

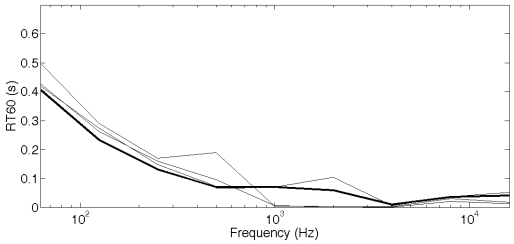

Figure 4: Measured $RT_{60,3D}$ (thick line) and corrected $RT_{60,2D}$ (thin lines) for a DWM with porous absorber.

5.5. Cubic room with pillar

The final example considered is the diffuse room described in Section 5.3 but with a pillar which runs from floor to ceiling and whose cross-sectional area is 1.83 m^2. The effect of this pillar is to increase surface area but decrease volume in the 3D case, and to increase the total boundary length but decrease the area in the 2D case. In both cases this has the effect of reducing the RT_{60}. The predicted and measured $RT_{60,3D}$ are 0.24s and 0.36 and the same $RT_{60,2D}$ values are 0.40 and 0.55.

Again, there are significant discrepancies between prediction and measurement, but those discrepancies are similar in both the 2D and 3D cases. The $RT_{60,2D}$ of the corrected 2D response is 0.27 which is closer to the $RT_{60,3D}$ than the un-corrected version, but is not as closely aligned compared to previous examples. Figure 5 shows the 3D, 2D and 3D-corrected 2D impulse responses and here an improvement in the corrected version can be clearly seen.

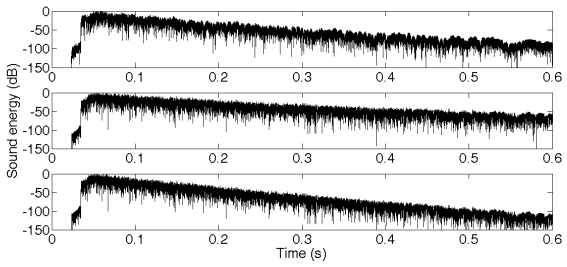

Figure 5: Energy plots of 3D (top), 2D (middle) and corrected 2D impulse responses for cubic diffuse room with pillar.

6. CONCLUSIONS

This paper has described, with reference to relevant theory, a method for matching the reverberation time of the outputs from 2D and 3D models of the same acoustic space for use in a hybrid room modelling system. A number of examples of simple rooms have been given and, in each case, the technique has produced an improvement in the correspondence between $RT_{60,2D}$ and $RT_{60,3D}$. Typical rooms that *RenderAIR* is designed to model are much more complex than this in terms of their geometry and the distribution of absorption and diffusion across surfaces and in terms of their spectrum. The simple rooms presented here are, due to their shape, highly modal and frequency-dependency has only been considered in one example. Readers who are interested in how different aspects of the *RenderAIR* system combine to produce accurate simulations of real-world spaces can find more information in [2]. A more detailed analysis of how this temporal alignment algorithm performs on real-world spaces will be the subject of future work.

7. REFERENCES

[1] M.R. Schroeder, "Natural sounding artificial reverberation", *Journal of the Audio Engineering Society*, October 1962, pp. 219-223

[2] D.T. Murphy et al., "Hybrid room impulse response synthesis in digital waveguide mesh based room acoustics simulation", *Proceedings of the 11th International Conference on Digital Audio Effects* (*DAFx-08*)

[3] A. Southern and D. Murphy, "Low complexity directional sound sources for finite difference time domain room acoustic models", *Paper presented at the 126th Audio Engineering Society Convention,* 2009.

[4] A. Southern and D. Murphy, "A second order differential microphone technique for spatially encoding virtual room acoustics", *Paper presented at the 124th Audio Engineering Society Convention,* preprint 7332, 2008.

[5] D.T. Murphy et al., "Acoustic modeling using the digital waveguide mesh", *IEEE Signal Processing Magazine*, March 2007.

[6] S. Bilbao, *Wave and Scattering Methods for Numerical Simulation*, John Wiley and Sons Ltd., Chichester, UK, 2004.

[7] E. Thompson, *The soundscape of modernity: architectural acoustics and the culture of listening in America, 1900-1933,* MIT Press, Cambridge, USA, 2002.

[8] C.F. Eyring, "Reverberation time in "dead" rooms", *Journal of the Acoustical Society of America*, 1930, pp. 217-241.

[9] C.W. Kosten, "The mean free path in room acoustics", *Journal of the Acoustical Society of America*, 1960, pp. 245-250.

[10] M. Karjalainen and C. Erkut, "Digital waveguides versus finite difference structures: equivalence and mixed modeling", *EURASIP Journal on Applied Signal Processing*, June 2004, pp. 978-989.

[11] K. Kowalcyzk and M. van Walstijn, "Modeling frequency-dependent boundaries as digital impedance filters in FDTD and K-DWM room acoustics simulations", *Journal of the Audio Engineering Society,* July/August 2008, pp. 569-583.

[12] K. Kowalcyzk and M. van Walstijn, Private communication, 2008.

[13] M. E. Delany and E. N. Bazley, "Acoustical Properties of Fibrous Absorbent Materials," *Journal of Applied Acoustics*, 1970, pp. 105–116.