REAL-TIME WALKTHROUGH AURALISATION OF THE ACOUSTICS OF CHRIST CHURCH CATHEDRAL, DUBLIN

G Kearney Department of Theatre, Film and Television, University of York, United Kingdom
M Gorzel Department of Electronic and Electrical Engineering, Trinity College Dublin, Ireland
F Boland Department of Electronic and Electrical Engineering, Trinity College Dublin, Ireland
F Smyth Urban Institute, University College Dublin, Ireland
D Lennon Urban Institute, University College Dublin, Ireland
H Rice Department of Mechanical and Manufacturing Engineering, Trinity College Dublin, Ireland

1 INTRODUCTION

Amongst Ireland's many acoustic treasures, Christ Church Cathedral, located in the heart of Dublin city, represents one of the most historically significant. The Cathedral has not only consistently played a prominent part in Dublin's civic as well as religious history but has also been noted throughout the ages for the calibre of its choir and music. A key element is the acoustics of the Cathedral, whose varying degrees of responsiveness and clarity are notable at different registers. Despite the fact that much research has been undertaken into the architectural and musical history of this cathedral, prior to this work, its acoustics have not been studied, nor measured for posterity. This is in fact the case for many of Dublin's important acoustic splendours, including St. Patrick's Cathedral, located only a short distance from Christ Church. The architectural mapping and measurement of the acoustics of such spaces is vital to Ireland's cultural heritage preservation.

A significant body of research has been undertaken in the measurement of the acoustics of historically important performance spaces for posterity, most notably in Italy[1]. Current methodologies involve the acoustic measurement of performance spaces using Soundfield microphones and binaural mannequins[2]. Anechoic recordings of musical performances can then be filtered (convolved) with these measurements giving a plausible formation of an auditory scene in a virtual version of the acoustic environment as it exists today. This is referred to as 'Data-based auralisation'.

Whilst such methods allow us to take a 'snapshot' of the acoustic, they do not accommodate a fully interactive listening experience, as would be the case in real-world listening. The acoustic changes that occur due to the movements of both the source and listener within the space are an important aspect in the perception of music performance[3]. However, current methods in computational-based auralisation correctly consider this aspect through computing grids of acoustic response measurements over the entire computer model for different source/receiver configurations[4,5]. Real-time auralisation is then implemented by choosing the appropriate acoustic response (or interpolating across responses[6]) for the immediate source/listener positions. However, grid measurements are unfeasible in the context of real-world acoustics as they are tedious and prone to error. Furthermore, one must not neglect the fact that performance within reverberant spaces such as cathedrals will be markedly different than that in an anechoic chamber, due to the performer interaction with the acoustic. In this regard, it is a challenge to be able to capture the performances within the reverberant space such that this aspect is preserved whilst recording for an interactive auditory scene.

In this paper, we demonstrate how real-world acoustic and architectural measurements in combination with computational-based auralisation can be used to create a plausible, interactive and real-time walkthrough auralisation and computer graphic-visualisation for the purpose of acoustical heritage preservation of Christ Church Cathedral. This paper is outlined as follows: In Section 2 we will outline the historical significance of Christ Church Cathedral as a performance space in Dublin City. In Section 3 we will then demonstrate our approach to acoustic measurement and performance recording within the space for the purpose of walkthrough auralisation. In Section

4 we show how architectural scans of the Cathedral can be utilised for interactive visualisation in a real-time gaming engine, and Section 5 then discusses the implementation of the audio engine. Finally, Sections 6 and 7 detail the public demonstration of the model and draw conclusions to the paper.

2 HISTORICAL BACKGROUND TO CHRIST CHURCH CATHEDRAL

The Romanesque origins of Christ Church date from c. 1030, and it was rebuilt and dedicated c. 1172 as an Augustinian Priory under Archbishop Laurence O'Toole. Christ Church became an Anglican cathedral in 1541[7]. Changing religious and social contexts within Dublin in the nineteenth century led to a decrease in the available funds for upkeep of the cathedral which inevitably fell into disrepair, and was subject to a rather controversial architectural restoration in the late nineteenth century undertaken by the architect George Edmund Street and funded privately by a philanthropic citizen named Henry Roe[7]. The controversy had a number of facets, including the fact that the cathedral was restored according to an earlier medieval footprint, and the plan altered from (roughly) cruciform to apsidal. A shortened quire and the inclusion of a choir screen were further causes of discontent, and the choir screen was slightly modified in 1883 after Street's death in an ostensible bid to improve the acoustics[8]. More recent restorative work, notably on the crypt and cathedral office has been carried out by Paul Arnold Architects.

A series of grants and stipends were made available at the end of the fifteenth century to enable the cathedral to school and support four choirboys[8]. The first of the grants, which came from a prominent citizen, carried further stipulation with regard to the development of polyphonic music within the Cathedral, in recognition of the manner in which such investment would elevate the status of the city of Dublin. Some 400 years after this, Henry Roes financial support in the nineteenth century made provision for the restoration of music as well as architectural fabric at Christ Church. This was with a bequest to provide for an increase in choral music within the cathedral as well as a permanent endowment of its choir (to be composed of twelve efficient choirmen, along with a minimum of twelve boy choristers and an organist)[9]. This investment in music is significant, as the built fabric of Christ Church in the early nineteenth century had suffered from lack of finances to the point where it required significant restorative and structural work in the latter half of that century, yet the salaries offered to musicians were quite substantial in comparison with similar positions in England. As such, the cathedral was able to attract very talented musicians, and thereby foster a strong and highly regarded tradition of liturgical music[8]. The publication by Street and Seymour in 1882[9], documenting the history and restoration of the cathedral, noted that

> *"Though Christ Church Cathedral, with its grand endowment for the choral services, never lost a certain sort of attraction for the people, it was wholly a musical one, and none ever seemed to trouble themselves to reflect on the beauties of the building, or on the evil state into which it had been allowed to fall"*
> G. E. Street and E Seymour. The Cathedral of the Holy Trinity commonly called Christ Church Cathedral Dublin: an account of the restoration of the fabric. London, Sutton, Sharpe and Co., 1882.

3 DATA ACQUISITION

3.1 Acoustic Measurement

The historic significance of Christ Church fully justifies the requirement to measure its acoustics as they currently exist for cultural heritage preservation. To this end, acoustic measurements were taken from two source positions. The first source position was to the centre of the choir, the second was in front of the choir where large-scale choral events are often staged. This was to provide a comprehensive assessment of both musical scenarios. An omni-directional dodecahedron loudspeaker, as specified by ISO-3382 was used for acoustic excitation[10]. The receivers used consisted of a Soundfield MK5 system and a Neumann KU100 binaural head as shown in Figure 1.

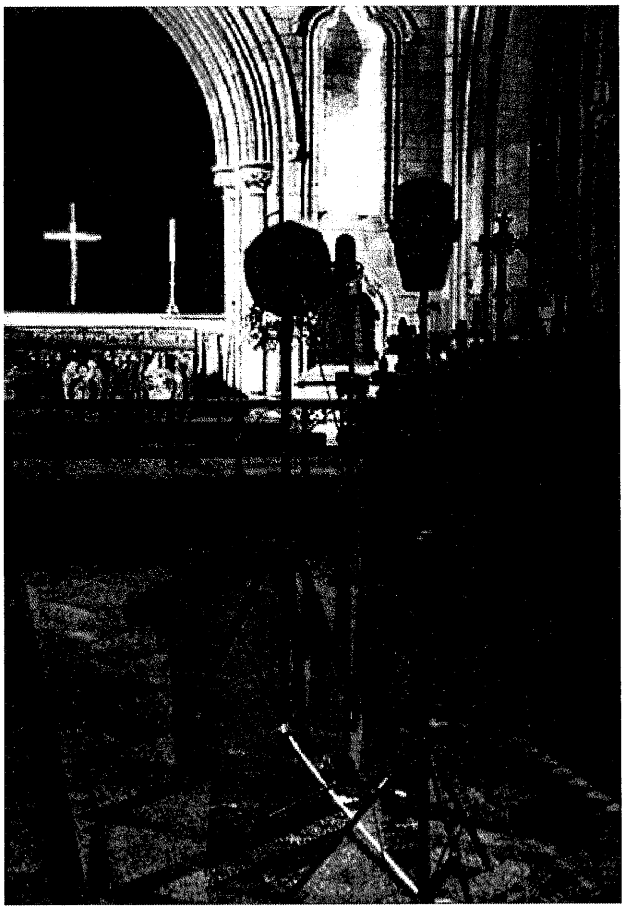

Figure 1. Acoustic measurements taken using omnidirectional source, soundfield and binaural microphones

Measurements were taken along the length of the nave, and to either side of the pews within the nave. Each measurement position was replicated for each source position. Measurements were taken from both source points at 1, 2, 4, 8,16, 24 and 32 metre increments from each source. The additional height of staging was accounted for in determining the height of the source position. The excitation signal used was a logarithmic sine swept tone of duration 60 seconds, ensuring that any speaker induced distortion in the resultant impulse responses was removed[11]. A detailed account of these measurements and results has recently been documented by the authors[12].

In this paper, we pay particular attention to the two major acoustic parameters of Reverberation Time (RT) and Inter-Aural Cross Correlation (IACC). The reverberation time is a measure of the decay of an impulsive sound source in a reverberant room, defined as the time it takes for the impulse to fall by 60dB after the direct sound. IACC gives us a measure of the similarity of the ear signals. In general, the higher the value of the IACC, the narrower the perceived source width. If this parameter is measured within the first 80ms of the impulse response, the influence of the early reflections can be ascertained. This parameter is specifically relevant in the 500, 1000 and 2000Hz bands, where the wavelengths involved are comparable to the dimensions of the head, and an average IACC of these bands is known as the $IACC_{E3}$[13].

The reverberation time for source position 1 (centre choir position) is shown in Figure 2. The resultant spatially averaged reverberation time at 1kHz was measured as 3.2 seconds. It can be observed that from 8m, the reverberation time does not change across all frequencies and the diffuse field properties are dominant. The changes in IACC for the same source-receiver positions,

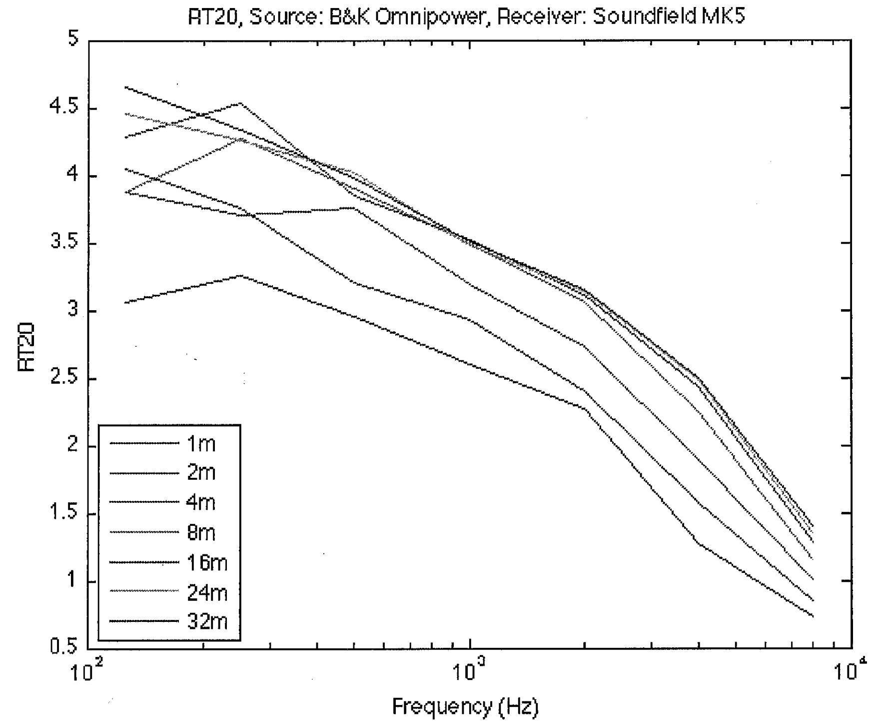

Figure 2. Reverberation time measurements measured at Nave receiver positions from source position 1

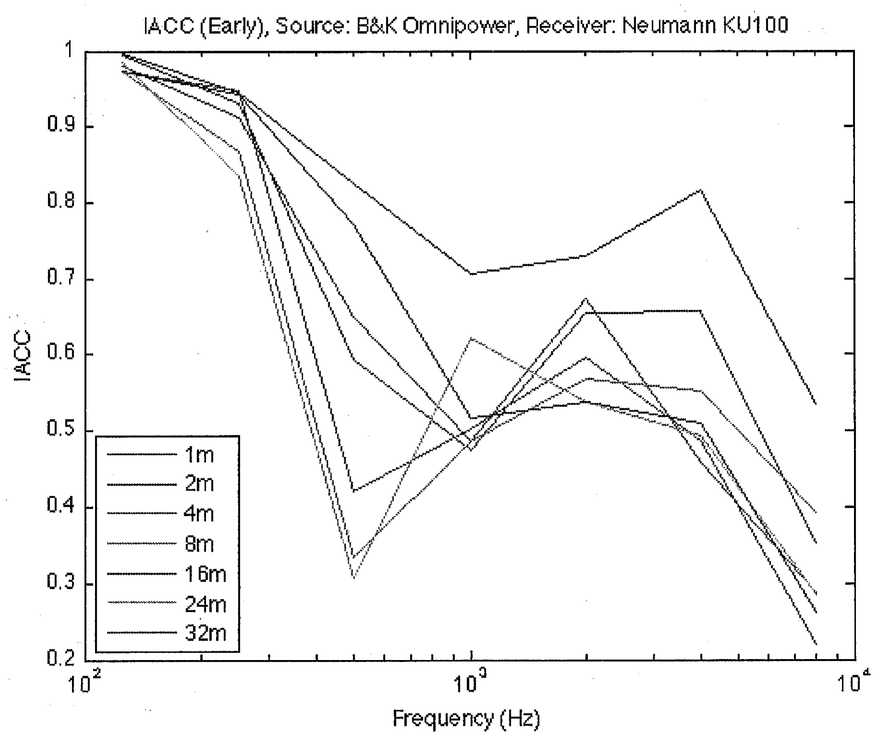

Figure 3. IACC measurements measured at Nave receiver positions from source position 1

are shown in Figure 3. In general, the IACC values beyond 1m are similar, with noticeable changes around 500Hz. The spatially averaged $IACC_{E3}$ is measured as 0.5717.

Finally, in-situ absorption coefficient measurements were taken of a number of key surfaces using the methodology of Mallais[14]. These were later used to inform the general acoustic reflective properties in the modelling process.

3.2 Choral Recording

The strong musical tradition associated at Christ Church has been upheld to the present day and the calibre of the cathedral choir is outstanding. In order to create an effective virtual model of the cathedral, it is therefore also necessary to capture a real performance as sung by the choir. An anechoic recording is insufficient, as the resultant performance is not a true representation of the choir interaction with the acoustic, i.e. the performance characteristics and dynamics would be significantly different. The choir was therefore recorded in the Cathedral during a Sunday service in June 2010.

Direct-field capture of each singer was achieved within thier critical distance (the distance from the source at which the reverberant field energy equals the direct sound energy) using a spot microphone. The positioning and directional characteristic of the microphone is important not only to the tonal balance, but also to minimize the amount of other sources (commonly referred to as spill) in the recorded signal. Uni-directional microphones are frequently used in order to maximize rejection, but the cost of increased directional response can often lead to compromised frequency response in lower grade microphones as well as proximity effect. Such frequency response distortions must therefore be corrected in post processing. Thus, each of the 19 choir singers was provided with the individual cardioid microphone (Rode NT5). Although microphones were placed well within the critical distance of each singer, it was impossible to avoid cross-talk from other singers. However, the signal to noise ratio was deemed acceptable (greater than 10dB at each microphone) and did not cause any major tonal distortions during the process of auralisation.

Reference recordings were also made within the Cathedral to later compare with the auralisation. These were made at the choir centre position using a tetrahedral microphone array, and at the crossing, using the Neumann KU100 binaural head and a set of AKG C-414 microphones configured as an ORTF pair.

(a) (b)

Figure 4. Choir recording (a) Spot microphones on each performer circled in red (b) Rode NT5 in custom mount with popshield.

3.3 Geometric and Spatial Data Capture

A comprehensive spatial and architectural survey of Christ Church Cathedral was carried out in order to build a detailed model with accurate surface, height, and volumetric information. To ensure both accuracy and speed in the data collection, the work was carried out using a terrestrial LiDAR (Light Detection and Ranging) laser scanner at a tightly set matrix to collect surface and height information and thus build up a three dimensional representation of the interior of the cathedral. The age and lengthy restoration history of the building have resulted in a number of volumetric and geometric idiosyncrasies. These constitute an essential part of the cathedrals character, but would be very difficult and time consuming to capture accurately using more traditional surveying techniques such as triangulation and macro-modelling. LiDAR captures information using laser pulses set to a user-specified matrix with distances calculated by measuring the time delay between emitted and reflected pulses, and is currently the swiftest way of capturing this type of information accurately[15].

The model used in surveying Christ Church Cathedral was a Leica C10. The C10 captured data in individual scans at 360° along the horizontal axis and 270° along the vertical axis, thereby allowing for the capture of awkward geometries with relative speed and minimal overlap. Six-inch semi-spherical targets were used to aid the scanning process, and their locations were registered to the scanner for recognition and subsequent automatic merging of individual scans. The entire process was completed in the space of 5 hours with a total of 16 individual scans taken from calculated scanner- and-target locations. Each of the 16 scans was carried out at a 10m by 10mm resolution at a range of 10m. The majority of the cathedral was captured in the first four scans. The geometry of the main body of the cathedral is quite open, but the interlocking spaces and built-up nature of the area around the transepts and the Lady Chapel required a greater density of scan stations and targets for complete data collection. A built-in, dual-axis level compensator to the scanner allowed for inconsistencies to floor level to be fully taken into account. The collected data from each scan station was saved directly to the hard drive of the scanner - an integrated data manager runs Windows XP - then downloaded to an external PC laptop for post-processing. The registered targets were then aligned and automatically merged using Leicas proprietary Cyclone software.

The laser-captured data sets are in point cloud format, which consists of millions of individual points. Data in this form is read by standard architectural drawing packages as a solid-block object. To this end, the survey data had to be converted into a workable spatial representation by translation into a series of planes and surfaces. This was carried out with the assistance of Severn Partnership, with the results exported to AutoCAD.

4 REAL-TIME VISUALISATION

4.1 Data Reduction

Since the purpose of this work was to present an audio-visual walkthrough demo of the Christ Church Cathedral choir performance in Dublin, it was necessary to faithfully model the interior space of the church. For this purpose we have chosen to visualise the model in Blender 3D 2.49b software[16]. Blender 3D is an open source, fully integrated 3D modelling/animation environment with a built-in game engine (BGE). It has a full support for the Python scripting language[17], which greatly extends its scope of possible applications and allows, for example, for real-time communication with other software (like Pure Data[18]). It also enables to import and work with professional architectural models created in AutoCAD (.dxf files).

The laser scanning process (LiDAR survey described earlier) allowed the formation of a detailed and highly accurate (high-poly) AutoCAD model of the church interior. Initially, after import into Blender 3D, the geometry consisted of 95580 polygons (triangles), 85674 vertices and was divided into 23 smaller, manageable sections. The initial wireframe model is shown in Figure 5. Although such a detailed model would be interesting to work with from the off-line analysis point of view, it was obvious that the level of complexity was significantly too high for real-time video rendering, not

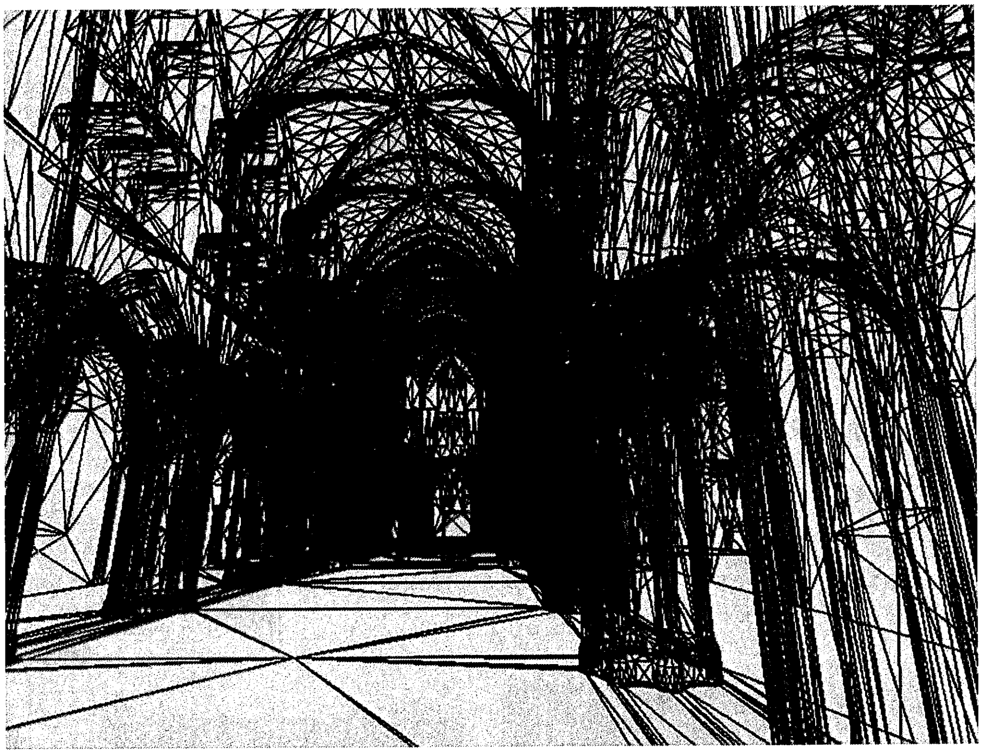

Figure 5. High poly version of the Christ Church model (wireframe)

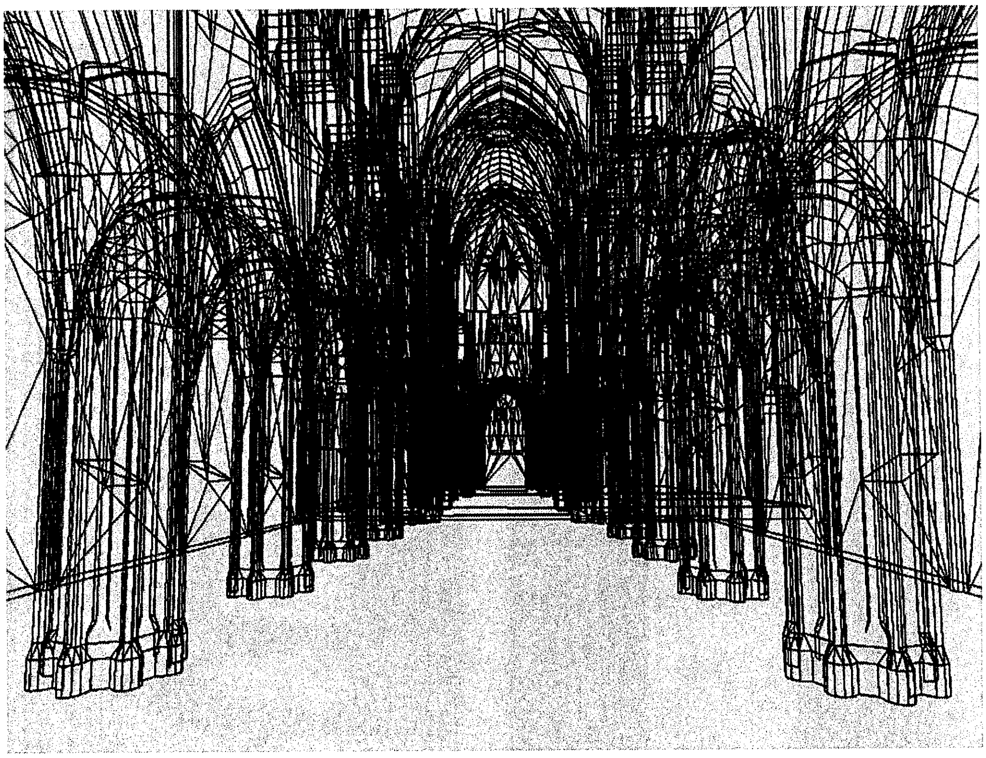

Figure 5. Poly-reduced quadified version of the Christ Church model (wireframe)

to mention even the simplest geometrical acoustic simulation methods. Despite the fact that modern PC games can use environments with hundreds of thousands or even millions of polygons we also have to remember that they use state-of-the-art commercial visual game engines and quite basic audio engines. Not to be forgotten is that they also require top quality hardware to show the full potential at acceptable frame rates (fps). On the other hand, we wanted our demo to run on an off-the-shelf PC.

Thus, we have employed optimisation process that allowed obtaining a simplified version of the mesh without perceptual loss of detail. The process consisted of removing doubled or nearly adjoining vertices, quadification (converting many triangular faces to single quads natively supported by BGE) and poly-reduction. For the visual presentation, the initial model proved to be highly redundant and after the optimisation process, we reduced the poly count to 35107 and the number of vertices to 30892, which was enough to run the demo on a mid-class PC (Dell Optiplex, 4-core Intel processor, 4 GB of RAM, ATI Radeon GPU with 256 MB of RAM). This reduction is shown in Figure 6.

The model was also textured and lit in order to achieve a more natural appearance. Textures were created mostly from the photographs taken in the interior of the church and some example textures are shown in Figure 7. The ambition here was not to obtain a full, faithful, photorealistic replica of the space but rather to improve the overall visual impression and recreate a feeling of spaciousness that was somehow distorted in a solid model.

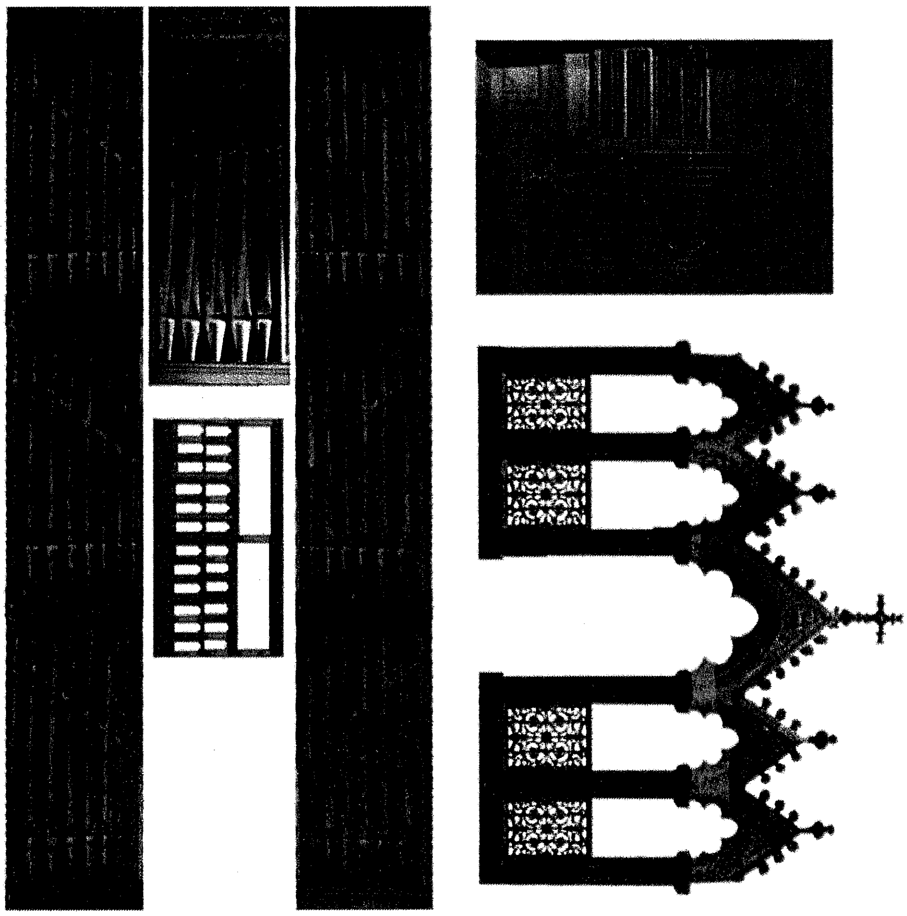

Figure 7. Example textures incorporated into model

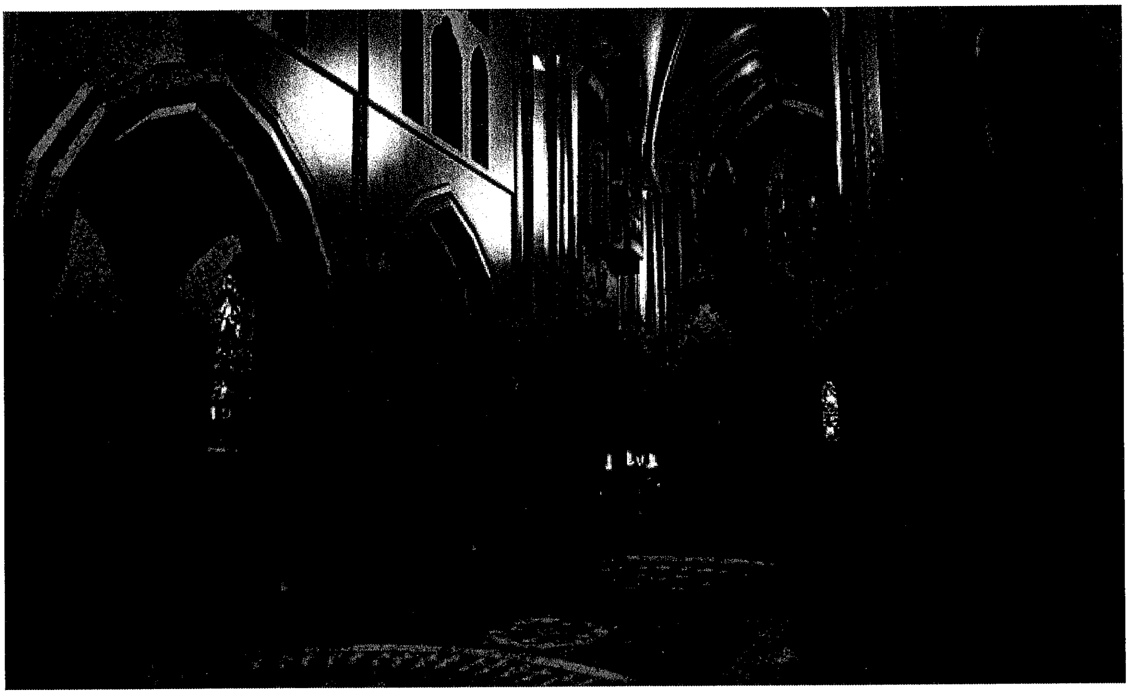

Figure 8. 1st person view of the current textured model

4.2 Adding Interactivity and Integration with the PD audio engine

Adding interactivity to the finished model comprised of adding a virtual camera to the model and making it responsive to PC controllers: mouse and keyboard. A simple python script was written to assign mouse movement to camera orientation and keyboard input to camera translation, which are typical settings in many video games. It made it possible for the user to move freely and explore the interior of the church model. To represent the sound sources in the 3D space, silhouettes of choir singers were also added as simple 2D textures that were programmed to always face the virtual camera. In the demo, coordinates of the virtual singers correspond to actual positions in the church where the recordings took place and are fixed throughout the whole demo. Parameters that do change in the interactive model are: the position and orientation of the camera in the virtual world, as well as horizontal angles between the camera and each of the sound sources. From the auralisation point of view, whenever the camera position/orientation changes, the listening perspective changes as well and needs to be updated in the audio engine. Thus, we have established a unidirectional communication between Blender 3D and Pure Data that allowed communicating the users position/orientation and source angle data at runtime. The communication was done using UDP protocol and was governed by another python script.

5 REAL-TIME AURALISATION

5.1 User Interaction

Having constructed a Virtual Visual Environment (VVE) of the cathedral, we now consider the creation of the virtual acoustic model. The objective is to present an aurally plausible reproduction of the captured performance in a virtual version of the cathedral, herein termed the Virtual Auditory Environment (VAE). An important aspect in the quest for realism in such auditory scene synthesis is user interaction. That is, how the movements of a person listening to the virtual auditory scene directly influence the scene presentation.

Such walkthrough auralisation presents several challenges for production engineers, the most significant of which is the generation of the correct room acoustic response due to a given source-listener position. In particular, the correct direction of arrival of the direct sound and early reflections must be maintained since these signals contain the most vital cues for localisation of acoustic sources.

In recent years, the formation of auditory scenes based on real world spaces has benefited greatly from the use of convolution reverberation techniques, and a significant body of work has been presented illustrating the possibilities and limitations[19, 20, 21]. The representation of room responses in this manner assumes that the source-room interaction is one that is linearly time-invariant (LTI). In reality, the room impulse response (RIR) changes significantly with the relative spatial positions of the source and the listener. However, if the geometric properties of the performance space are known (as well as the frequency dependent absorptive properties of the materials in the room), then the acoustic response can instead be computed. Highly accurate results can be achieved using wave-based methods, such as the Finite Element Method (FEM), Boundary Element Method (BEM)[22] or Finite Difference Time Domain (FDTD) method[23]. Due to computational expense, such methods are generally limited to low frequency RIR estimation. Geometric-based solutions to calculating RIRs, such as the image-source method[24] are well suited to the mid-to high frequency regions, although they do not consider phenomena such as diffraction or scattering. However, calculation of the propagation delays and magnitudes at low reflection orders using image sources is well suited to real-time auralisation. Hybrid reverberation algorithms have been proposed which combine computational and measured impulse responses but have largely focused on the synthesis of the diffuse decay as opposed to early reflections[25,26]. In this paper, we focus on the real-time rendering of the early reflections in conjunction with pre-rendered diffuse field recordings for walkthrough auralisation.

5.2 Audio Engine Implementation

The audio engine used for auralisation was fully implemented in the Pure Data visual programming environment[18]. In the ideal scenario, a full-length impulse response should be obtained and applied for each of the sound sources in the space at runtime. Since the real-world impulse response measurements can easily reach millions of samples, their realisation as FIR filters is computationally expensive. For this reason, we implemented the basic methodology previously described by the authors[27] which utilises a hybrid model of the RIR by decomposing it into two parts: the short deterministic part (direct sound + early reflections) that is synthesised at runtime and the long, pre-recorded and pre-convolved diffuse reverberation. The deterministic part can be calculated using one of the geometrical approaches (like Image Source Method) whereas stochastic reverb is obtained from the real Spatial Room Impulse Response captured from a soundfield microphone. The diffuse part of the impulse response is obtained from the directional analysis and diffuseness estimation of Merimaa and Pulkki[28].

In our simplified model, sound sources (singers) are represented as 2D textured planes that are always facing the virtual camera and each of the planes is constantly casting a ray in its normal direction. The ray-casting algorithm returns true whenever the virtual camera is visible to the sound source and false whenever there is an obstacle on the ray's path (like a wall or a pillar). In the case of baffles with semi-transparent textures (like fences) they hold a special property flag that makes them invisible to the algorithm. The visibility flag (true/false) is subsequently passed to the audio engine and determines whether to calculate or mute the direct sound from a particular sound source. The ray-casting algorithm is one of the basic BGE built-in logic modules.

The Image Source method was chosen to synthesise early reflections parts of new impulse responses. However, despite apparent simplicity, the ray tracing based methods in general become quickly computationally expensive when complex geometries are used since the number of virtual sound sources tends to grow rapidly with added faces. Visibility checking and calculating image

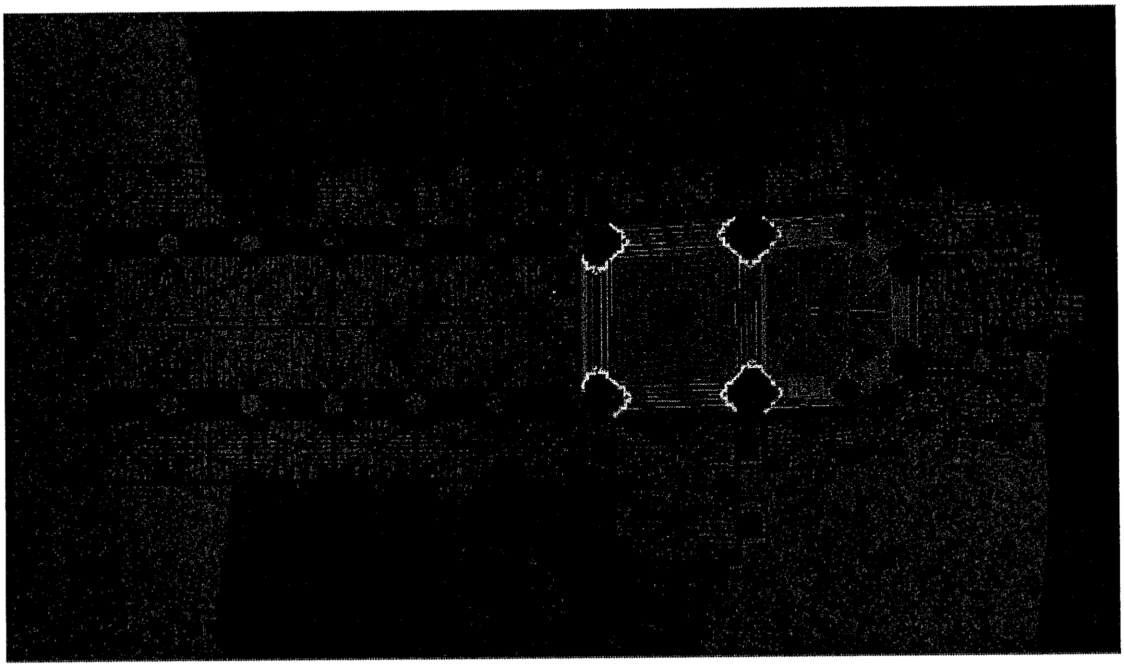

Figure 9. Bounding box of the highlighted part is used to calculate image sources whenever user enters the choir area.

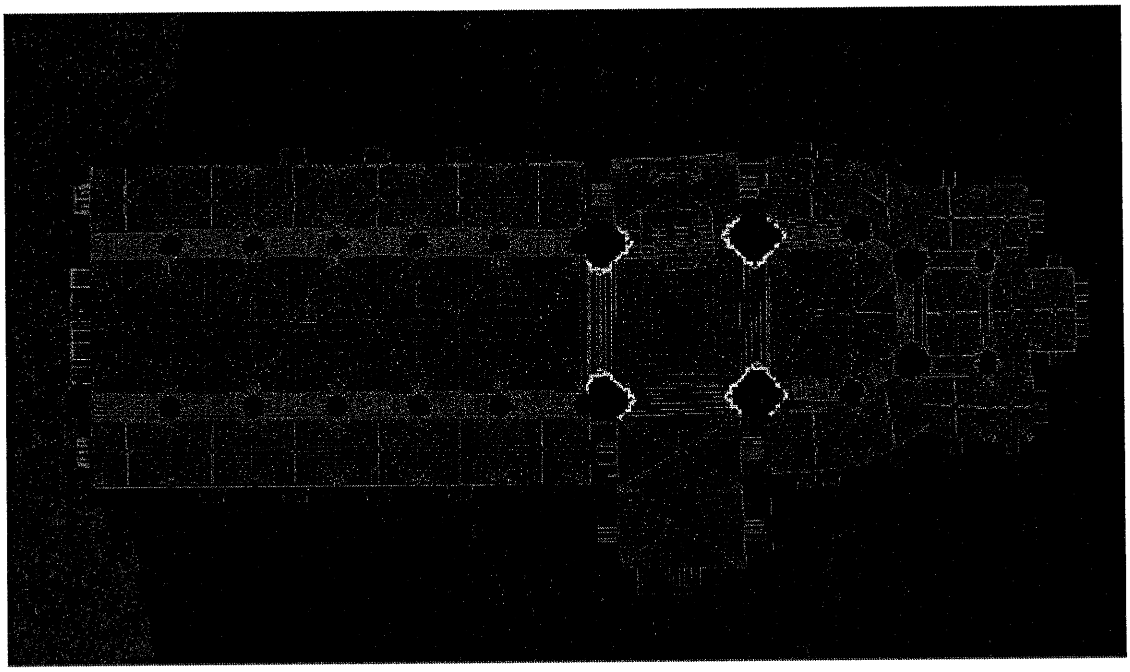

Figure 10. Bounding box of the whole model is used to calculate image sources whenever user leaves the choir area.

sources for a >35k poly model (and also handling as many channels of audio) would not be feasible considering a current state-of-the-art technology so simplifications were unavoidable. As a first approximation for calculating image sources in our model we use two bounding boxes of the interior depending on whether the user is inside (Figure 9) or outside (Figure 10) the choir area. The global bounding box has the 2D dimensions of 60 x 24 meters whereas the inner bounding-box is 20 x 11 meters. The audio engine is then set to logically switch between the two sets of mirror images whenever the user enters/leaves the choir area. For each sound source we create a map of image sources using the method described by McGovern[29]. The method allows for fast calculation of image sources that are stored in regular matrices (3x3, 5x5 etc. for 2D case or 3x3x3, 5x5x5 etc. for 3D case) where the number of elements of the matrix 1 indicate the total number of image sources. If the virtual sources do not change their locations at runtime it is possible to pre-calculate the image sources and store them in lookup tables.

In our case we represent each sound source as a direct sound and eight reflections (3x3 matrix), which is realised in Pure Data as a multi-tap variable delay network. Delay times are recalculated based on the actual distance from the receiver to the direct/image source. Low-pass filtering (smoothing) is applied to delay curves in order to avoid zipper noise. Intensity values of each of the sources obey the inverse square law and we flip the phase of the signal upon each reflection in order to avoid unipolar responses with strong DC offset. In this way, we represent each sound source as nine audio streams giving 171 streams in total when we combine together all the sources in the scene. For direct source/receiver distances greater than 8m, image sources are not calculated and the auralisation is only of the direct sound and diffuse field. This is due to the fact that at this distance there is little perceptible difference with or without the early reflections in the model, which is also corroborated by the largely consistent acoustic measurements of Section 3 beyond 8m.

5.3 Spatialisation

The reproduction of the auditory scene is also highly dependent on the spatialisation method employed. Many techniques have been proposed in the literature, most notably Vector Based Amplitude Panning (VBAP)[30] and Wavefield Synthesis[31]. However, Ambisonics[32], which is based on the spherical harmonic decomposition of the soundfield, represents a practical and asymptotically holographic approach to real-time soundfield manipulation. Ambisonics was originally developed by Gerzon, Barton and Fellgett[33] as unified system for the recording, reproduction and transmission of surround sound. The theory of Ambisonics is based on the decomposition of the soundfield measured at single point in space into spherical harmonic functions, and overviews of the theory and decoding methods can be found in the literature[3, 33].

In the next step, each audio stream is encoded into 3rd order Ambisonics taking into account the angle between the direct-source or image-source and the receiver (virtual camera). By combining the corresponding Ambisonic signals we create a spherical harmonic representation of the whole auditory scene. At this stage we also add the reverberation tail, which is obtained by pre-convolving a mono mix of all the sources with the diffuse component of the measured (1st order) soundfield impulse response. The level of reverberation can be adjusted using a slider and Wet/Dry ratio can be set to achieve the desired perceptual result.

In the last, decoding stage, we have a flexibility of choosing the final loudspeaker setup for spatial audio reproduction. As the default we've chosen the regular octagonal rig (horizontal only). Binaural rendering was also implemented using the Virtual loudspeaker approach[35], where Head-Related Impulse Responses (HRIRs) are measured at the sweet-spot (the limited region in the centre of a reproduction array where an adequate spatial impression is generally guaranteed) in a multi-loudspeaker reproduction setup, and the resultant binaural playback is formed from the convolution of the loudspeaker feeds with the virtual loudspeakers. In order to avoid the latencies resulted from multiple HRIR filtering we utilise the novel approach proposed by the authors allowing for significantly shorter HRIR runtime filters[36,37]. Using the Ambisonics approach here means that instead of recalculating the angle of arrival for every sound source and every reflection it is much easier to rotate the whole sound field according to the incoming head-tracking data. The headtracker used for demonstration was the InertiaCube2+ Inertial Measurement Unit (IMU) providing stable and accurate orientation data on three axes (3 degrees of freedom).

6 DEMONSTRATION

The current system was demonstrated to the public as part of National Heritage week at the Mansion House Dublin, August 2010. Whilst the current implementation requires further optimization and validation, informal listening comparisons between the real and virtual environments resulted in a markedly positive public response. Moreover, the project succeeded in raising public awareness about the importance of acoustic measurement of Irish historical spaces for cultural heritage preservation. The walkthrough model is also demonstrated at the poster stand of this event.

7 CONCLUSION

In this paper we have presented an interactive audio-visual model of Christ Church Cathedral, Dublin. The model is based on real-world acoustic and geometric measurements of the space. Hybrid reverberation was utilized, where early reflection synthesis was achieved using image source modelling, combined with the real-world diffuse field measurement. We have shown how real-world recordings of the choir could be achieved using direct-field pickup and subsequently auralised in a walkthrough implementation. Further work is required into the optimisation of the acoustic model, in particular the early reflection synthesis. A perceptual analysis is also required to verify the model as well as further verification with ISO-3382 acoustic parameters. The project has yielded a significant amount of information about the architectural heritage and the current record of Christ Church cathedral and has provided a solid framework for expanding the scope of such investigations within Ireland. Both the data collected, and the methodologies utilised are unique to Irish heritage.

8 ACKNOWLEDGEMENTS

This project has been funded, supported and facilitated by two Architectural Research Grants (2009 and 2010) from the Irish Heritage Council, Trinity College Dublin, University College Dublin, The Board and Chapter of Christ Church Cathedral Dublin, Leica and SIS Ireland and the Severn Partnership.

9 REFERENCES

1. Angelo Farina. Acoustic quality of theatres: Correlations between experimental measures and subjective evaluations. Journal of Applied Acoustics, 62:889–916, 2001.

2. Angelo Farina and Regev Ayalon. Recording concert hall acoustics for posterity. In 24th International Conference of the Audio Engineering Society, 2003.

3. Gavin Kearney. Auditory Scene Synthesis using Virtual Acoustic Recording and Reproduction. PhD Thesis, Trinity College Dublin, 2010.

4. Jens Huger Rindel, Claus Lynge, Graham Naylor, and Kim Rishoj. The use of a digital audio mainframe for room acoustical auralization. In 96th Convention of the Audio Engineering Society, 1994.

5. B.-I. Dalenbäck and M. Strömberg. Real time walkthrough auralization - the first year. In Proceedings - Institute of Acoustics, vol. 28, Pt.2, 2006.

6. G. Kearney C. Masterson and F. Boland. Acoustic impulse response interpolation using Dynamic Time Warping. In 35th International Conference of the Audio Engineering Society, February 2009.

7. E. K. MILNE. Christ Church Cathedral, Dublin: A History. Four Courts Press, 2000.

8. R. Stalley. George Edmund Street and the Restoration of the Cathedral, 1968 - 78. Christ Church Cathedral, Dublin: A History, Milne, K. (Ed.), Four Courts Press, 2000.

9. G. E. Street and E Seymour. The Cathedral of the Holy Trinity commonly called Christ Church Cathedral Dublin: an account of the restoration of the fabric. London, Sutton, Sharpe and Co., 1882.

10. ISO. ISO 3382-1:2009 acoustics - measurement of room acoustic parameters - part 1: Performance spaces, 2009. http://www.iso.org/, accessed 21-09-2009.

11. A. Farina. Simultaneous measurement of impulse response and distortion with a swept-sine technique. In 108th Convention of the Audio Engineering Society, 2000.

12. F Smyth, D. Lennon, J. Ollie, G. Kearney, F. Boland, and H Rice. A Comparative Study of the Acoustics of Christ Church and St. Patrick's Cathedrals. Research Report 16857. Kilkenny, The Irish Heritage Council, 2009.

13. T. Okano, L. L. Beranek, and T. Hidaka. Relations among interaural cross-correlation coefficient (IACC), lateral fraction (LF), and apparent source width (ASW) in concert halls. The Journal of the Acoustical Society of America, 104(1): 255–265, 1998.

14. Scott Mallais. In situ determination of acoustic absorption coefficients. In 125th Convention of the Audio Engineering Society, 2008.

15. F. Smyth and D. Lennon. Modelling and retro-modelling: Lidar and point cloud as a tool in acoustic analysis. Urban Institute Ireland Working Paper Series, 09, 2009.

16. Blender Foundation. Blender. http://www.blender.org/, accessed 20th November, 2010.

17. Python Software Foundation. Python programming language. http://www.python.org/, accessed 20th November, 2010.

18. Miller Puckette. Pure data, 2009. http://puredata. info/, accessed 20th November, 2010.

19. Ralph Kessler. An optimised method for capturing multidimensional acoustic fingerprints. In 118th Convention of the Audio Engineering Society, May 2005.

20. V. Pulkki and J Merimaa. Spatial Impulse Response Rendering I: Analysis and Synthesis. Journal of the Audio Engineering Society, 53, 2005.

21. G. Kearney and J. Levison. Actual Vs. Virtual multichannel acoustic recording. In 134th Convention of the Audio Engineering Society, May 2008.

22. Andrzej Pietrzyk. Computer modelling of the sound field in small rooms. In 15th International Conference of the Audio Engineering Society, 1998.

23. L. Savioja, T. Rinne, and T. Takala. Simulation of room acoustics with a 3-D finite difference mesh. In International Computer Music Conference, pages 463–466, 1994.

24. J. B. Allen and D. A. Berkley. Image method for efficiently simulating small-room acoustics. JASA, 65: 943–950, 1979.

25. R. Stewart and D. Murphy. A hybrid artificial reverberation algorithm. In Proc. of the 122nd AES

Convention, Vienna, Austria, 2007.

26. S. Browne. Hybrid reverberation algorithm using truncated impulse response convolution and recursive filtering, 2001. http://www.music.miami.edu/programs/mue/mue2003/research/sbrowne/, University of Miami, FL, USA., accessed 20th November, 2010.

27. M. Gorzel, G. Kearney, H. Rice, and F. Boland. Virtual acoustic recording: An interactive approach. In 10th International Conference on Digital Audio Effects (DAFX), September 2010.

28. J. Merimaa and V. Pulkki. Spatial impulse response rendering I: Analysis and synthesis. Journal of the Audio Engineering Society, 53, 2005.

29. Stephen G. McGovern. Fast image method for impulse response calculations of box-shaped rooms. Applied Acoustics, 70(1): 182-189, 2009.

30. V. Pulkki. Virtual sound source positioning using Vector Base Amplitude Panning. Journal of the Audio Engineering Society, 45:456-466, 1997.

31. A. J. Berkhout. A Holographic Approach to Acoustic Control. Journal of the Audio Engineering Society, 36:977-995, 1988.

32. M. A. Gerzon. Practical periphony: The reproduction of full sphere sound. In 65th Convention of the Audio Engineering Society, February 1980.

33. M. A. Gerzon. Periphony: With-height sound reproduction. Journal of the Audio Engineering Society, 21:2–10, 1973.

34. Bruce Wiggins. An investigation into the real-time manipulation and control of three dimensional soundfields. Ph.D Thesis, University of Derby, 2004.

35. A. McKeag and D. McGrath. Sound field format to binaural decoder with head-tracking. In 6th Australian Regional Convention of the AES, 1996.

36. G. Kearney, C. Masterson, S. Adams, and F. Boland. Approximation of binaural room impulse responses. In IET Irish Signals and Systems Conference, June 2009.

37. Gavin Kearney, Claire Masterson, Marcin Gorzel, Henry Rice, and Frank Boland. Application of HRIR Factorization to Game Audio. In 41st International Conference of the Audio Engineering Society, 2011.