Portal-based Sound Propagation for First-Person Computer Games

Cameron Foale
University of Ballarat
Mt. Helen, Australia
c.foale@ballarat.edu.au

Peter Vamplew
University of Ballarat
Mt. Helen, Australia
p.vamplew@ballarat.edu.au

ABSTRACT
First-person computer games are a popular modern video game genre. A new method is proposed, the *Directional Propagation Cache*, that takes advantage of the very common portal spatial subdivision method to accelerate environmental acoustics simulation for first-person games, by caching sound propagation information between portals.

1. INTRODUCTION
Sound is vitally important in creating a sense of player immersion in first-person computer games. Sound contributes strongly to the sense of presence, or of "being in the game" [43, 6]. Sound cues are used by players to determine the in-game locations of non-visible events and objects, and to reinforce visual information, helping players to gauge the distance of visible objects [37]. Further, sound is an important dimension for artistic input in the design of games.

Acoustics auralization is the process of simulating how a sound would be heard if it was placed in an environment. Auralization requires calculating how a sound propagates from a sound source to a listener. Sound propagation from a sound source to a listener is described by an *acoustic transfer function*.

Existing games use relatively unsophisticated acoustic transfer functions, such as simple low-pass filters and generic delay and reverb effects. These functions are computationally cheap. However, the resulting sound environment is only loosely based on how sound would actually propagate around the environment.

Valve's Half-life 2 [42] presents a more sophisticated approach, interpolating the parameters of delay and reverb effects between artist-specified settings, or *soundscapes*, based on the current environment. For example, a sound designer creates a large reverberant room soundscape, and a small reverberant room soundscape, and the audio engine interpolates between these depending on the size of the room containing the player. This method can produce highly stylised artistic effects, however there is no attempt to model sound propagation, and the effects do not depend on the direction and location of the sound source or listener.

The most common acoustic transfer function in architectural acoustics is the *impulse response*. An impulse response describes how a single "spike" of sound pressure (an *impulse*) at the source position arrives, over time, at the listener position. The simplest way to conceptualise the impulse response is to make a loud, sharp noise (such as a hand clap or bursting balloon) in a real environment - the combination of echoes heard is an impulse response, unique to the source location and the position of the listener. A recorded or generated impulse response can be mathematically applied to any sound signal using convolution, to "place" the sound at the location that the impulse was created, however convolution is a computationally expensive process without dedicated hardware.

Simulating an accurate impulse response in real-time is too computationally expensive for computer games. However, for other aspects of computer games, performing calculations off-line and then applying the results in real-time is common, such as offline calculation of global illumination offline for real-time lighting and shading [20, 25], and offline spatial subdivision and potentially visible set calculations for real-time rendering [1]. We suggest that there is a large demand for higher quality sound simulation in computer games, and there is scope to use off-line computations to improve real-time auralization.

1.1 Structure
This paper is structured as follows:

- Sections 2-5 provide background on auralization and sound phenomena
- Sections 6-8 describe existing autralization methods
- Sections 9 and 10 describe the new method, and means for evaluation

1.2 Assumptions
The following simplifying assumptions and approximations are common in architectural acoustics literature, and are also made for this paper:

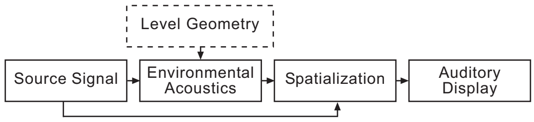

Figure 1: Conceptual Audio Pipeline

1. The geometry of the environment is largely static. This assumption allows us to consider off-line calculation of wave propagation.
2. The speed of sound is constant throughout the environment, simplifying propagation calculations.
3. Objects in the environment all move subsonically, and usually much slower. This allows us to avoid considering non-linear wave phenomena such as shockwaves and sonic booms, which can be emulated in computer games outside of the wave propagation simulation.

2. AUDIO PROCESSING PIPELINE

Figure 1 shows a generic audio processing pipeline that is useful for discussion purposes, similar to that described in [38]. *Source Signal* is the dry, anechoic signal such as the sound of gunfire or speech; *Environmental Acoustics* is the process by which the effects of sound wave propagation are applied to the Source Signal, informed by Level or World Geometry; *Spatialization* is the process by which source direction is simulated in multiple sound channels; and the *Auditory Display* is the physical output to speakers or headphones. This process occurs in parallel for each environmental sound source, with the final mixing stage only occuring at the auditory display stage. Non-environment sounds such as virtual radio communications or control feedback may skip the environmental acoustics and spatialization stages, and we do not further consider these types of sounds.

Note that sound signals generated by rigid body dynamics simulations, scripting and voice communications have been excluded from Figure 1, but may be considered part of the Source Signal generation.

3. SOUND PERCEPTION

The way the human ear perceives the direction and distance of a sound can be described using several binaural characteristics; absolute amplitude, the Inter-aural Amplitude Difference (IAD), the Inter-aural Time Delay (ITD), subtle filtering by the ear pinnae, and environmental effects.

The first three effects are handled by the *spatialization* system in a sound engine, the last by the environmental acoustics simulation, combined with spatialization of reflections.

Environmental effects contribute largely to the sense of presence in the environment [43], and the sense of the size of the environment. However, environmental effects do not necessarily increase spatial localization accuracy. Lateral reflections decrease localization accuracy for simple environments [18, 29], particularly when early lateral reflections are relatively high in amplitude compared to the direct sound. However, coincident or nearly-coincident reflections increase localization accuracy [18], and little research has been done regarding localization in complex environments.

Sounds are commonly spatialized binaurally (with headphones or in-ear phones) or using a set of discrete speakers.

4. ENVIRONMENTAL ACOUSTICS

4.1 Propagation

A *point source* in a uniform medium will generate a spherical wave. Ignoring air absorption, the energy of the wave is conserved as it propagates. For a spherical pressure wave, the energy is distributed evenly across the spherical surface area given by $4\pi r^2$[35]. As the wave propagates and r expands, the energy per unit area decreases following the inverse square law, $1/r^2$. The energy of a wave is proportional to the square of its amplitude, $E \propto A^2$, and so the amplitude decreases as $1/r$. At $22°$Celcius, the speed of sound in air at sea level is 343ms^{-1}.

When an acoustic wave travels from one medium to another, some of its energy is *reflected*, some is *absorbed* (transformed into another form of energy), and some is *transmitted* into the second medium. Most commonly, for environmental acoustics purposes, the concern is with the interaction of a wave travelling through air and interacting with solid materials, although other fields such as ultrasound and sonar are concerned with effects between other media.

4.2 Absorption

When an acoustic wave crosses a boundary between two materials, some energy is absorbed. The absorption properties of a material are frequency and angle dependent, but independent of amplitude. For simulation purposes, it is useful to determine a coefficient or set of coefficients that describe the level of absorption, which can then be included directly into the simulation algorithm.

4.3 Transmission

Transmission occurs when an acoustic wave travelling in one medium crosses a boundary to another medium of different acoustic impedance. In these cases, energy is transmitted in the form of an acoustic wave travelling in the second medium. In the common case in architectural acoustics and FPS computer games, sound sources and listeners are rarely located in media other than air, and so the noticeable effects of transmission are those when sound travels from air, to another medium, and back to air again. Depending on the medium through which the wave travels the effect may be hardly noticeable, such as sound heard through a piece of fabric, to a marked reduction in high frequencies, such as sound heard through a masonry wall.

The study of sound transmission is generally with a view to improving the sound insulation properties of materials and rooms. Sound Transmission Loss (STL) is measured as the amount of energy lost, in decibels, on transmission through a material between coupled rooms. Similarly to absorption, STL is frequency but not amplitude dependent [44], and is therefore a linear time-invariant process which can be modelled using digital filters or impulse responses.

4.4 Reflection

Reflection occurs when an acoustic wave travelling from one medium to another is partially reflected according to Snell's law. The amplitude and phase of the reflection depend on the absorption properties of the medium, and the directivity of the reflection depends on the surface structure.

Reflection can be characterised as *specular*, *diffuse*, or somewhere in between. Specular reflection follows the simple law of reflection, $\theta_i = \theta_r$; the angle of reflection is equal to the angle of incidence. This essentially means that for a specularly reflecting surface, no energy is *scattered* upon reflection. Specular reflection occurs when the reflecting surface is smooth relative to the wavelength, and thus a surface can exhibit both specular and diffuse reflection depending on the frequency of the incident wave.

In contrast, diffuse reflection occurs when the surface is rough relative to wavelength. Incoming wave energy is scattered in many directions; depending on surface features and uniformity of the absorption coefficient, the energy may be mostly redirected back towards the incident wave, scattered in all directions equally, or scattered in a more complex fashion.

4.5 Diffraction

Diffraction is the scattering of waves by obstacles of size similar to the wavelength. Diffraction allows waves to “bend” around corners and obstacles. The amount a diffracted wave will bend depends on the wave frequency. Lower frequencies diffract more dramatically than higher frequencies.

Diffraction has been largely ignored in real-time architectural acoustics, due largely to the geometric models used (see section 6.2). However, it is widely accepted that simulation of diffraction effects are important to obtain correct results [7, 40, 39]. The development of the Uniform Theory of Diffraction (UTD) - a geometric description of wave diffraction - has enabled geometric models to incorporate diffraction effects.

4.6 Reverberation

Reverberation is caused by the recursive reflection and diffraction of a wave in an acoustically “live” environment - that is, an environment where at least some surfaces are reflective.

Early studies into reverberation were concerned with either predicting simple RT_{60} reverberation time[10, 33] or room absorption characteristics [28, 11, 27].

The reverberation characteristics of a reverberant path from a source to a receiver for an environment (with no non-linear resonances, such as rattling window panes) is a linear time-invariant (LTI) filter, and as such can be captured by its *impulse response* [17, 13].

5. IMPULSE RESPONSES

An impulse response is the time-domain response of an LTI system to a single *impulse*, or “spike” of sound.

By capturing or creating an impulse response, it is then possible to simulate the application of the LTI filter to a signal by *convolution*. Note that not all LTI filters have finite impulse response lengths, however for acoustics use we assume that impulse responses decay to the limit of floating-point accuracy with no continuous oscillation.

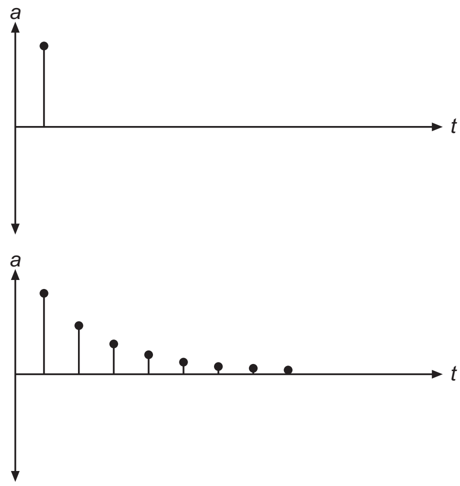

Figure 2: *Top:* A single impulse, and *Bottom:* An impulse response representing the output of an exponentially decaying feedback delay

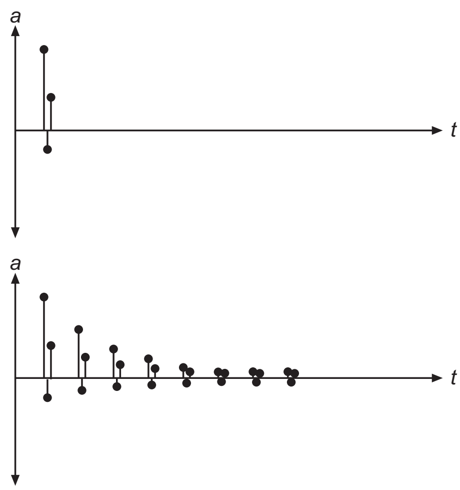

Figure 3: *Top:* A “dry” signal, and *Bottom:* The resulting signal after convolution with the impulse response in figure 2

6. SIMULATING ACOUSTICS

The computer simulation of acoustic environments has been an active area of research for several decades. The fundamental problem for computer simulation is determining the audible direct and indirect sound contributions due to reflection and diffraction for a given source and receiver pair. While this is currently possible using offline computation, reverberation time is commonly in the order of several seconds, and direct calculation of all possible propagation paths is currently impossible for real-time simulation.

[14] broadly group the existing methods into those that model sound propagation as a wave phenomenon and those that model propagation as a geometric (ray-based) phenomenon. We will follow this taxonomy in the following discussion, and refer the reader to [14] for a more complete review.

6.1 Wave Models

6.1.1 Finite Element Method and Boundary Element Method

Finite Element Methods (FEM) and Boundary Element Methods (BEM) solutions directly solve the wave equation, either through a spatial partitioning scheme in the case of FEM, or through a surface partitioning scheme for BEM. These methods can produce accurate models of wave travel and diffraction, however computational cost rises exponentially with the highest representable frequency. FEM and BEM approaches naturally incorporate diffraction effects[41], however the computational requirements have previously limited their use in real-time applications. In architectural acoustics, BEM and FEM methods are generally only used for modelling diffraction effects below 150 Hz [41]. For non-trivial environments, BEM and diffuse raytracing yield largely equivalent results[5]. Given the simpler implementation of raytracing over BEM methods, we do not further consider the BEM.

6.1.2 Finite Difference Schemes and Digital Waveguide Networks

Finite Difference Scheme (FDS) methods discretize space similarly to the FEM approach. FDS uses the *finite difference approximation* to the first- and second-order derivatives of the pressure term in the wave equation. By discretizing on a mesh (commonly rectilinear), with uniform time steps, the wave amplitude at each point can be computed using only amplitude information from the previous two time steps; this approach is called the Finite Difference Time Domain (FDTD) scheme.

Digital waveguide networks (DWNs) model wave propagation as a network of nodes connected by bi-directional single-sample delays [30, 35], and for rectilinear grids DWNs yield equivalent results to FDTD methods[34]. The performance of both DWNs and FDS methods is too slow for real-time evaluation at audio frequencies, however they are suitable for precomputing and caching impulse responses on a per node level [31], and have been used successfully for modelling low-frequency sound propagation [26, 36].

6.2 Geometric Models

Geometric models treat sound propagation as a linear, ray phenomenon. That is, geometric methods largely ignore the transmission medium other than to fix a constant wave speed, and the modelling of propagation is done purely by geometric operations.

6.2.1 Image Source Method

The first of the geometric methods to be investigated for computer simulation of acoustics was the image-source method [3], based on Eyring's formulation of reverberation [10]. In this approach, specular sound reflection is modeled using "virtual" sources, created by reflecting the sound source over every polygon visible to it in the mesh representing the environment. Further specular reflections are modeled by recursively repeating the process. To generate a final sound, an impulse response is calculated by combining the delays and reflection properties of every image "visible" to the receiver. This method has exponential complexity for all but the simplest contrived environments, and after very few reflections the visibility check becomes prohibitively expensive.

6.2.2 Ray Tracing

The ray tracing method can be used in acoustics modelling, and is used very similarly to its application in computer graphics[24]. A number of infinitely thin rays are traced from the sound source, reflecting from walls and diffracting at corners. Those rays that eventually intersect with the sound receiver (or a sphere surrounding the receiver) contribute to the impulse response describing the acoustic transfer. This form of ray tracing has problems very similar to light-to-eye ray-tracing in computer graphics - there may be large sampling error (even large features may be missed by all rays), and in the common case the majority of rays traced never intersect with the listener.

It is difficult to treat diffraction correctly using ray tracing, as the infinitely thin rays will only rarely intersect with infinitely thin edges. Further, to determine the final contribution of all rays, the listener must be represented as a volumetric element, an approximation that may lead to simulation errors. However, ray tracing has the major advantages over other simulation methods of intuitiveness and ease of implementation.

Ray-tracing is an inherently scalable algorithm - the number of reflections, the number of rays cast and the number of diffuse rays generated upon reflection can be adjusted dynamically, however ray tracing is typically too computationally intensive to attain results of reasonable quality in real-time.

Several extensions to ray tracing have been used in acoustics, including Cone Tracing [4] and Pyramid Tracing [12]. These methods reduce, to some degree, the sampling errors that arise from tracing a finite number of rays.

6.2.3 Beam Tracing

Beam tracing methods [19, 8, 9] trace volumetric elements called *beams* through the environment, which are reflected using geometric clipping and mirroring operations. For audio purposes, beam tracing has mainly been used as a preprocessing step to construct a *beam tree*, which is used to accelerate a subsequent ray tracing step [14].

Beam tracing helps avoid the sampling issues associated with distributed ray tracing, however the geometric calculations involved are complex enough that deep beam trees must be computed off-line, which largely limits the algorithm to fixed geometry, and fixed sources or receiver.

6.3 Statistical Reverberation

Statistical reverberation methods are built around capturing the perceptible essense of an acoustic environment, rather than the exact impulse response. The term *statistical reverberation* stems from the observation that, for high enough frequencies, both the phase and frequency response of late reflections are essentially randomly distributed, regardless of the shape of the room [32, 16]. Statistical reverberation algorithms, or *digital reverberators*, use the randomness of late responses to reduce the level of computation required to simulate long reverberations.

For a detailed history of the development of digital reverberators, see [16].

Digital reverberators are extremely common for use in computer games, as they are much cheaper computationally than convolution with large impulse responses. Further, many consumer sound cards provide hardware acceleration for digital reverberation effects. Digital reverberators as used in current games do not spatialize late reflections even though this is perceptually important [43]. Reverberator parameters can be created to match calculated impulse responses [22], however this approach is not suitable for acoustic environments that exhibit non-exponential reverberation decay [12].

7. SPATIALIZATION

7.1 Binaural Reproduction

Binaural reproduction attempts to simulate how sound arrives at the ear, including IAD, ITD and pinnae effects. While the ITD in particular is easy to implement using simple delay lines, the frequency and phase modifications caused by pinnae are more complex.

For a given direction at each ear, a Head Related Transfer Function (HRTF) can be measured [15] or calculated [23, 2] to describe the filtering effects due to pinnae, head and torso diffraction. This function can then be applied to source signals, and fed to each ear individually using headphones.

Unfortunately, HRTFs calculated for one person are not necessarily usable by other people. Anthropometric differences such as head size, hair, pinnae shape and body size all contribute to the effect of a single HRTF, and so these functions are typically usable in their raw impulse-response state by a narrow group of users. For consumer hardware, generic HRTFs generated by analysis of averaged recordings are used for binaural spatialization.

7.2 Surround Reproduction

Surround reproduction involves playing sounds to be spatialized through one or more of a set of discrete speakers that "surround" the listener.

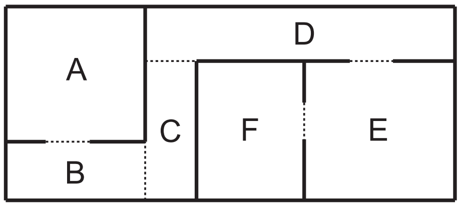

Figure 4: Geometry can be divided into *cells* (A, B, C, D, E, F) and interconnecting *portals* (AB, BC, CD, DE, EF)

7.3 Use of Impulse Responses

Many of the above wave and geometric methods for simulating environmental acoustics are built around the idea of constructing an impulse response per speaker, in a combined acoustics simulation and spatialization stage. This method results in the most accurate transfer of the simulation results to the auditory display, however the final impulse responses depends upon the locations and orientations of the source and listener. If the source or listener moves, the entire simulation must be re-run.

This problem does not occur for methods in which the environmental simulation is decoupled from spatialization, such as using a simple digital reverberator. In this case, only the direct contribution must be re-spatialized.

8. FIRST-PERSON COMPUTER GAMES

Current generation first-person computer games perform relatively limited acoustic environment modelling. Most game engines simply combine spatialization of direct sound with a single non-spatialized digital reverberator for environmental simulation. The parameters available for adjusting the acoustics simulation are those available to the reverberator, and may be limited to simply switching between preset parameter values. The direct sound may be low-pass filtered to simulate diffraction and transmission effects. The volume of a sound source is determined by artist-specified maximum and minimum amplitude distances, with no relationship to physical sound propagation.

9. PROPOSED METHOD

A new technique is proposed, the *directional propagation cache*, to cache offline sound propagation calculations.

The caching mechanism takes advantage of the *portal subdivision* method, in which areas of space (*cells*) are separated by invisible partitioning polygons, called *portals*. Portal subdivision has been used in prior works to accelerate computer graphics algorithms, for both off-line and real-time applications. Portal methods have also been used in first-person computer games to accelerate visibility, audibility and artificial intelligence operations, see for example [21]. There are several methods of generating a portal-based spatial decomposition. The most common are using the Binary Space Partition (BSP) method, and specifying portals manually.

The technique stores, for a directed sound wave travelling through each portal, the contribution and directivity of the reflected waves arriving at every other portal. We call the cache of all directed responses for a single pair of portals a

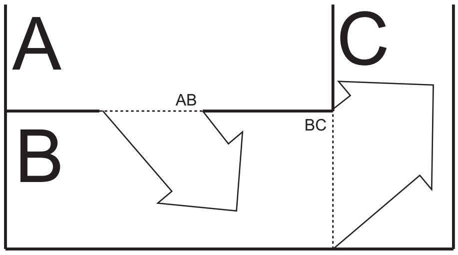

Figure 5: Can the way sound travels from portal AB to portal BC be used in real-time?

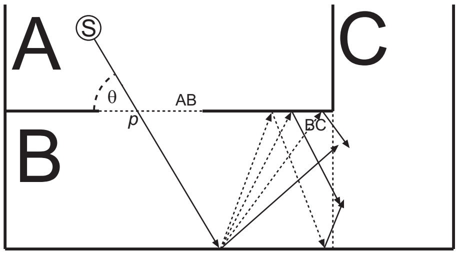

Figure 6: A portal response $PR(AB{\rightarrow}BC)$can be viewed as, for each possible ray crossing portal AB, the collection of reflected, diffracted and transmitted rays crossing portal BC

portal-response, and the collection of portal-responses as a *directional propagation cache*. Note that a portal-response may use any acoustic transfer model, such as impulse-responses, digital filters or digital reverberator parameters.

9.1 Propagation Calculation

The core idea in the proposed method is to pre-calculate the sound propagation between portals, and use this information to approximate full propagation calculation in real-time (see Figure 5). Following the geometric theory of wave propagation, as used in ray-tracing, we can describe the response between portals AB and BC as a collection of rays travelling through AB and the associated set of rays arriving at BC (see figure 6). Note that as the portal response mechanism is independent of the propagation calculation, we call the response derived from ray-tracing a *ray-portal response*.

We denote a ray leaving portal AB as $P_{(AB,p,\theta)} \rightarrow$, where θ is the angle of incidence to the portal, and p is the ray-portal intersection point. We denote a ray entering portal AB as $\rightarrow P_{(AB,p,\theta)}$. The ray-portal response is then defined as:

$$PR_{(AB\rightarrow CD,u,v,\theta,\phi)} = \sum P_{(AB,u,\theta)} \rightarrow P_{(CD,v,\phi)} \qquad (1)$$

where AB and CD are portal identifiers, $P_{(x)} \rightarrow P_{(y)}$ denotes a ray travelling from x to y, u and θ are the point and angle at which a ray exits portal AB, and v and ϕ are the point and angle at which reflected rays enter portal CD. Note that we are specifically describing the ray-portal response in terms of rays rather than sound pressure or intensity. This allows us to see that the the response from one portal to another depends on 4 variables of 2 dimensions each, in 3-dimensional space (see Figure 6). Here then is an opportunity to take advantage of spatial coherence and perceptual limitations, by grouping rays into buckets based on their angular variables. We can further take advantage of spatial coherence by discarding the tracking of u and v, the portal intersection points.

By adopting this grouping method, we can discard the maintainance of individual ray angular information and simply store the other ray variables, such as travel time, amplitude and spectral information. This leads to the definition of a ray-portal response of equation 2.

$$PR_{(AB\rightarrow CD,i,j)} = \sum P_{(AB,i)} \rightarrow P_{(CD,j)} \qquad (2)$$

where i,j are ray group identifiers.

Finally, we combine the contribution of all rays for a given angular group into a single impulse response for direct convolution, or for analysis for other transfer function models. No further offline calculations are necessary.

9.2 Run-time

Run-time behaviour consists of four distinct steps:

1. Determine the contribution of the sound source to each portal in the source cell (source portals)
2. Determine the contribution of each portal in the listener's cell (listener portals) to the listener
3. Perform a cache lookup to find the transfer function from each source portal to each listener portal
4. Create point sound sources on or near listener portals for spatializing the final response

Determining contributions to and from source and listener portals depends largely on the acoustic transfer function selected. The simplest method is to measure the angles formed between a vector from the source or listener position and the nearest point on the portal, and the portal surface itself. This angle can be used to select a single set of transfer function parameters from the cache, or to interpolate between several cached parameter sets.

Point sound sources are generated and placed on or near the portals in the cell containing the listener. We call these sources *imposter sources*, as they are placed to imitate the direction from which the propagated sound arrives. This method separates environmental acoustics simulation from spatialization, allowing the listener to change orientation without recomputing the environmental acoustics component. Further, generating a set of point sources can reduce computational cost by taking advantage of spatialization algorithms in audio hardware.

9.2.1 Propagation Between Distant Cells

This is perhaps the simplest case to conceptualize. For sound travelling from one cell to another, and where no direct signal progation is possible, the portal responses contain the complete set of information required to simulate spatialized propagation.

9.2.2 Propagation Between Nearby Cells

For nearby cells, the algorithm is very similar to the above case for distant cells, with the main difference being that direct sound may also be audible. Note that the offline simulation method does not consider the contribution of rays to portal responses until they have been reflected at least once. This allows the direct sound to be spatialized directly in parallel with the simulation algorithm, without including the same propagation path twice.

9.2.3 Internal Cell Propagation

For internal cell propagation, we must consider sound that leaves the cell, and re-enters after transformation. However, by treating portals as single sided, the above algorithms naturally handle sound leaving and re-entering cells with no special treatment.

9.2.4 Boundary Conditions

How to handle conditions when the listener (or source) crosses a portal boundary is somewhat more complex than it appears, and there are several potential solutions. The simplest is to perform an audio crossfade when the boundary is crossed. This has the advantage of handling any discontinuities, however the length of the crossfade may become quite long if the listener moves slowly across the portal border. Handling of boundary conditions depends largely on the transfer function and the placement of imposter sources.

10. EVALUATION

The proposed auralization system must be evaluated from both system quality and performance perspectives. A direct comparison of results from the directional propagation cache method with recorded impulse responses from a real environment is very difficult except in certain contrived cases [40], requiring aquisition of material absorption and diffusion parameters accurate enough to obtain physically correct results. Comparing the results to a ray traced solution would be simpler, as both algorithms can share the same geometry and material properties. In this case, we can quantitatively compare reverberation times or frequency spectra, or subjectively compare simulation quality.

However, comparing the proposed method to expensive offline simulations is not appropriate to the goals of the research. We wish to improve the quality of information presented by auralization in first-person games within specific performance parameters, so the most direct comparison is with the auralization in existing games. To this end, it is anticipated that user trials will be performed to determine whether players are better able to make judgements the virtual space with the directional propagation cache system.

Users will be asked to determine the location of a virtual sound source in a complex environment, with both the proposed method and a typical game audio engine, given foreknowledge of the structure of the environment. An improvement in the success rate and time taken to find the source would then show that the user is gaining more information about their virtual surroundings. This task-based approach is similar to other measures in the VR literature.

System performance will be measured by processor and memory utilization. These measurements are simple to perform and correspond well to subsystem performance measurements commonly used in the game industry [45].

11. REFERENCES

[1] J. M. Airey, J. H. Rohlf, and J. Frederick P. Brooks. Towards image realism with interactive update rates in complex virtual building environments. In *SI3D '90: Proceedings of the 1990 symposium on Interactive 3D graphics*, pages 41–50, New York, NY, USA, 1990. ACM Press.

[2] V. R. Algazi, R. O. Duda, R. Duraiswami, N. A. Gumerov, and Z. Tang. Approximating the head-related transfer function using simple geometric models of the head and torso. *The Journal of the Acoustical Society of America*, 112(5):2053–2064, 2002.

[3] J. Allen and D. Berkley. Image method for efficiently simulating small room acoustics. *Journal of the Acoustical Society of America*, 65(4):943–950, 1979.

[4] J. Amanatides. Ray tracing with cones. In *SIGGRAPH '84: Proceedings of the 11th annual conference on Computer graphics and interactive techniques*, pages 129–135, New York, NY, USA, 1984. ACM Press.

[5] A. L. Bot and A. Bocquillet. Comparison of an integral equation on energy and the ray-tracing technique in room acoustics. *The Journal of the Acoustical Society of America*, 108(4):1732–1740, 2000.

[6] E. Brown and P. Cairns. A grounded investigation of game immersion. In *CHI '04: CHI '04 extended abstracts on Human factors in computing systems*, pages 1297–1300, New York, NY, USA, 2004. ACM Press.

[7] P. T. Calamia and U. P. Svensson. Fast time-domain edge-diffraction calculations for interactive acoustic simulations. *EURASIP Journal on Advances in Signal Processing*, 2007:Article ID 63560, 10 pages, 2007. doi:10.1155/2007/63560.

[8] N. Dadoun, D. G. Kirkpatrick, and J. P. Walsh. The geometry of beam tracing. In *Proceedings of the Symposium on Computational Geometry*, pages 55–61, June 1985.

[9] I. A. Drumm and Y. W. Lam. The adaptive beam-tracing algorithm. *Journal of the Acoustical Society of America*, 107(3):1405–1412, 2000.

[10] C. F. Eyring. Reverberation time in "dead" rooms. *Journal of the Acoustical Society of America*, 1:217–241, 1930.

[11] C. F. Eyring. Methods of calculating the average coefficient of sound absorption. *The Journal of the Acoustical Society of America*, 4(3):178–192, 1933.

[12] A. Farina. *Pyramid Tracing vs. Ray Tracing for the simulation of sound propagation in large rooms*, pages 109–116. Computational Mechanics Publications, 1995.

[13] A. Farina. Room impulse responses as temporal and spatial filters. In *The 9th Western Pacific Acoustics Conference*, Seoul, Korea, June 2006.

[14] T. Funkhouser, N. Tsingos, I. Carlborn, G. Elko, M. Sondhi, J. E. West, G. Pingali, P. Min, and A. Ngan. A beam tracing method for interactive architectural acoustics. *Journal of the Acoustical Society of America*, 115:739–756, 2004.

[15] W. Gardner and K. Martin. HRTF measurements of a KEMAR dummy-head microphone. Technical Report 280, M.I.T. Media Lab Perceptual Computing, May 1994.

[16] W. G. Gardner. *Applications of Digital Signal Processing to Audio and Acoustics*, chapter 3, page 85. Kluwer Academic Publishers, 1998.

[17] M. A. Gerzon. Recording concert hall acoustics for posterity. *Journal of the Audio Engineering Society*, 23:569–571, September 1975.

[18] W. Hartmann. Localization of sound in rooms. *Journal of the Acoustical Society of America*, 74(5):1380–1391, Nov. 1983.

[19] P. S. Heckbert and P. Hanrahan. Beam tracing polygonal objects. *Computer Graphics*, 18, 1984.

[20] B. Hook. The quake 3 arena rendering architecture. In *Game Developer's Conference*, 1999.

[21] id Software. Doom 3. http://www.idsoftware.com, August 2004.

[22] J.-M. Jot. An analysis/synthesis approach to real-time artificial reverberation. *icassp*, 2:221–224, 1992.

[23] B. F. G. Katz. Boundary element method calculation of individual head-related transfer function. I. rigid model calculation. *The Journal of the Acoustical Society of America*, 110(5):2440–2448, 2001.

[24] U. Krockstadt. Calculating the acoustical room response by the use of a ray tracing technique. *J. Sound and Vibrations*, 8, 1968.

[25] J. Mitchell, G. McTaggart, and C. Green. Shading in valve's source engine. In *SIGGRAPH '06: ACM SIGGRAPH 2006 Courses*, pages 129–142, New York, NY, USA, 2006. ACM Press.

[26] D. T. Murphy, M. J. Beeson, J. Mullen, and D. M. Howard. Roomweaver and digitract: Two digital waveguide mesh acoustic modeling reseach tools. *The Journal of the Acoustical Society of America*, 116(4):2579–2579, 2004.

[27] R. F. Norris and C. A. Andree. An instrumental method of reverberation measurement. *The Journal of the Acoustical Society of America*, 1(1):32–32, 1929.

[28] P. E. Sabine. The measurement of sound absorption coefficients by the reverberation method. *The Journal of the Acoustical Society of America*, 1(1):27–27, 1929.

[29] S. Santarelli, N. Kopco, B. G. Shinn-Cunningham, and D. Brungart. Near-field localization in echoic rooms. *The Journal of the Acoustical Society of America*, 105(2):1024–1024, 1999.

[30] L. Savioja. *Modeling Techniques for Virtual Acoustics.* Doctoral thesis, Helsinki University of Technology, Telecommunications Software and Multimedia Laboratory, 1999.

[31] B. Schiettecatte, A. Nackaerts, and B. D. Moor. Real-time acoustics simulation using mesh-tracing. In *Proceedings of the International Computer Music Conference, Singapore*, 2003.

[32] M. Schroeder. Natural sounding artificial reverberation. *Journal of the Audio Engineering Society*, 10(3):219–223, 1962.

[33] R. S. Shankland. Architectural acoustics in america to 1930. *The Journal of the Acoustical Society of America*, 61(2):250–254, 1977.

[34] J. O. Smith III. On the equivalence of the digital waveguide and finite difference time domain schemes. http://arXiv.org/abs/physics/0407032, December 2005. Online Publication.

[35] J. O. Smith III. *Physical Audio Signal Processing: for Virtual Musical Instruments and Digital Audio Effects.* `http://ccrma.stanford.edu/~jos/pasp/`, 2006.

[36] J. O. Smith III and D. Rocchesso. Aspects of digital waveguide networks for acoustic modeling applications. Web published at http://ccrma.stanford.edu/ jos/wgj/., December 1997.

[37] J. Sodnik, S. Tomazic, R. Grasset, A. Duenser, and M. Billinghurst. Spatial sound localization in an augmented reality environment. In *OZCHI '06: Proceedings of the 20th conference of the computer-human interaction special interest group (CHISIG) of Australia on Computer-human interaction: design: activities, artefacts and environments*, pages 111–118, New York, NY, USA, 2006. ACM Press.

[38] T. Takala and J. Hahn. Sound rendering. In *SIGGRAPH '92: Proceedings of the 19th annual conference on Computer graphics and interactive techniques*, pages 211–220, New York, NY, USA, 1992. ACM Press.

[39] R. R. Torres, U. P. Svensson, and M. Kleiner. Computation of edge diffraction for more accurate room acoustics auralization. *Journal of the Acoustical Society of America*, 109(2):600–610, 2001.

[40] N. Tsingos, I. Carlbom, G. Elbo, R. Kubli, and T. Funkhouser. Validating acoustical simulations in bell labs box. *IEEE Computer Graphics and Applications*, 22(4):28–37, 2002.

[41] N. Tsingos, T. Funkhouser, A. Ngan, and I. Carlbom. Modeling acoustics in virtual environments using the uniform theory of diffraction. In *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 545–552, New York, NY, USA, 2001. ACM Press.

[42] Valve Corporation. Half-life 2. http://www.valvesoftware.com, 2004.

[43] A. Wakuda, H. Furuya, K. Fujimoto, K. Isogai, and K. Anai. Effects of arrival direction of late sound on listener envelopment. *Acoustical Science and Technology*, 24,4:181–185, 2003.

[44] R. Waterhouse. On standard methods of measurement in architectural acoustics. *Journal of the Acoustical Society of America*, 29:544–546, 1957.

[45] S. Woodcock. Game AI: State of the Industry. http://www.gamasutra.com, 1999.