



HAL
open science

Validation of Acoustical Simulations in the "Bell Labs Box"

Nicolas Tsingos, Ingrid Carlbom, Gary Elko, Thomas Funkhouser, Robert Kubli

► **To cite this version:**

Nicolas Tsingos, Ingrid Carlbom, Gary Elko, Thomas Funkhouser, Robert Kubli. Validation of Acoustical Simulations in the "Bell Labs Box". IEEE Computer Graphics and Applications, 2002, 22 (4), pp.28-37. 10.1109/MCG.2002.1016696 . inria-00606718

HAL Id: inria-00606718

<https://inria.hal.science/inria-00606718>

Submitted on 26 Jul 2011

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Validation of Acoustical Simulations in the “Bell Labs Box”

Nicolas Tsingos¹, Ingrid Carlbom¹, Gary Elko², Thomas Funkhouser³ and Robert Kubli²

¹Bell Laboratories - Lucent Technologies *

²MH Acoustics †

³Princeton University ‡

1 Introduction

Computer simulated sound propagation through 3D environments is important in many applications, including computer-aided design, training, and virtual reality. In many cases, the accuracy of the acoustical simulation is critical to the success of the application. For example, in concert hall and factory design (where OSHA sound limits must be met), the accuracy of the simulation may save costly re-engineering after construction. In virtual environments, experiments have shown that more accurate acoustic modeling provides a stronger sense of presence [13]. Furthermore, auditory cues aid in formation of spatial impressions, separation of simultaneous sound signals, and localization of objects [5], such as when a soldier locates an enemy in a training exercise or a fire-fighter locates a person stranded in a burning building. In contrast, incorrect auditory cues can lead to negative training [35, 6].

Although several systems (e.g., [10, 28, 9]) are available for computing sound propagation in 3D environments, there has not been detailed evaluation of their accuracy. The primary reason is that the acoustics of most real world environments is very complex, and thus detailed quantitative comparison of measured with computed impulse responses is difficult. As a result, acousticians have resorted to comparing gross statistics of impulse responses (e.g., reverberation time, clarity) [23] and/or using human listening tests for validation of computer simulations [1].




Figure 1: Left: a schematic view of the Bell Labs Box with measurement apparatus. Right: a computer graphics rendering of the Bell Labs Box lit by two point light sources.

This is similar to the situation in computer graphics during the mid-1980s. At that time, there were several global illumination algorithms that simulated the propagation of light through a 3D environment and produced an image for a given camera (e.g., ray

*The first author is now with the REVES research project at INRIA Sophia-Antipolis, France.

Nicolas.Tsingos@sophia.inria.fr

carlbom@research.bell-labs.com

†Summit, NJ, USA. gwe@ieee.org

‡funk@cs.princeton.edu

tracing [36] and/or radiosity [16]). Yet, there were very few results regarding the accuracy of the simulations. Computer-generated pictures and corresponding photographs rarely matched, and little insight into the causes of the mismatch could be derived from the comparisons due to the complexity of light transport in most scenes.

In response to this situation, researchers at Cornell University [16, 12, 25] constructed a simple real-world scene, the “Cornell Box”, in which the lighting, geometry, and reflectance properties of every surface were carefully measured and duplicated in a 3D computer graphics model. Using this simple and controlled experimental setup, they were able to make meaningful comparisons between global illumination simulations and photographs of the scene and provide explanations for the differences.

Motivated by the success of the Cornell Box, we have built a simple experimental setup for validating sound propagation simulations. We call it the “Bell Labs Box.” The setup comprises a simple configuration of planar surfaces with a speaker and a microphone (see Figure 1). While the basic configuration of the room is a simple six-sided box, it is constructed with reconfigurable panels that can be inserted or removed to create a variety of interesting geometries, including ones with diffracting panels (see Figure 2).

The key ideas behind this experimental setup are simplicity and control. First, the room has just a few planar surfaces, and thus possible sequences of reflections and diffractions are easily computed in a simulation and recognized in an impulse response. Second, the speaker radiation pattern, microphone directivity pattern, and reflectance properties of every surface are independently measured in an anechoic chamber, providing simulation parameters that closely match the real world environment. Finally, the room is configurable, providing a mechanism by which acoustical simulations can be validated with a variety of geometric effects. These features enable us to compare the results of simulations with measurement data in a simple and controlled setting, thereby allowing more detailed quantitative validation than has previously been possible.

2 Background and Previous Work

The problem addressed by this work is the validation of a solution to an integral equation expressing the wavefield (acoustic pressure and particle velocity) at some point in space in terms of the wavefield at other points (or equivalently on surrounding surfaces). For light simulations, the wave equation is described by Kajiyama’s rendering equation [18]. For sound simulations, it is described by the Helmholtz-Kirchoff integral theorem [8], which incorporates time and phase dependencies.

Since sound and light are both wave phenomena, the methods for their simulation and validation share many features. Yet, sound has characteristics different from light which introduce new and interesting problems:

- **Wavelength:** the wavelengths of audible sound are five to seven orders of magnitude longer than visible light, ranging between 0.02 and 17 meters (for 20KHz and 20Hz, respectively). Diffraction of sound occurs around obstacles of the



Figure 2: Additional panels can be mounted inside the Bell Labs box to study the effects of sound diffraction.

same size as the wavelength (such as tables) and reflections are primarily specular for large, flat surfaces (such as walls), while small objects (like coffee mugs) have little effect on the sound field (for all but the highest wavelengths). As a result, when compared to computer graphics, acoustics simulations tend to use 3D models with far less geometric detail. But, they must find propagation paths with specular reflections *and* diffractions efficiently, and they must consider the effects for different obstacles at a range of wavelengths.

- **Speed:** at 343 meters per second, the speed of sound in air is six orders of magnitude less than light, and sound propagation delays are perceptible to humans. Thus, acoustic models must compute the exact time/frequency distribution of the propagation paths, and the source sound must be auralized by convolution with the corresponding *impulse response*. This digital filter represents the delay and amplitude of the sound arriving along different propagation paths. In contrast, the propagation delay of light can be ignored and only the energy steady-state response must be computed.
- **Coherence:** sound is a coherent wave phenomenon, and interference between out-of-phase waves can be significant. Accordingly, acoustical simulations must consider phase when summing the cumulative contribution of many propagation paths to a receiver. More specifically, since the phase of the wave traveling along each propagation path is determined by the path length, acoustical models must compute accurate path lengths (up to a small percentage of the wavelength). In contrast, most light sources (except lasers) emit largely incoherent waves, and thus lighting simulations simply sum the power of different propagation paths.
- **Dynamic range:** the human ear is sensitive to five orders of magnitude difference in sound amplitude and arrival time differences allow some high-order reflections to be audible [37, 5]. Therefore, as compared to computer graphics, acoustical simulations usually aim to compute several times more reflections, and the statistical time/frequency effects of late sound reverberation are much more significant than for global illumination.

Despite these differences, the problems of simulation and validation are similar for both sound and light. For *simulation*, the main

difficulty arises from the wavefield discontinuities caused by occlusions, caustics, and specular highlights, resulting in large variations over small portions of the integration domain (i.e., surfaces and/or directions). Due to these discontinuities, no general-purpose, analytic formula can describe the wavefield at a given point, and solutions must rely upon sampling or subdivision of the integration domain into components that can be solved efficiently and accurately. Traditionally, four approaches have been used to address this problem: finite or boundary element methods [16, 26, 11], recursive ray tracing [22, 36], Monte Carlo path tracing [18, 28], and beam tracing [10, 24, 14, 17]. All four methods have been used for both sound and light.

For *validation*, the common problems are: (1) defining a quantitative measure of accuracy, and (2) understanding the causes of simulation errors. The first problem is generally addressed in acoustics with statistical measures of impulse responses and with human listening tests. The second problem has largely been explained by conjecture. For example, recent “round robin” studies applied several computer simulation tools to the same concert hall model and compared their output to *in-situ* measurements [34, 7]. Although, these studies provide a nice reference for comparing simulations, they differ from our work in several respects. First, they focus on *complicated concert halls*, rather than simple rooms like ours, and thus it is difficult to evaluate the correctness of any of the simulations (i.e., the acoustics are so complex that none of the tools get the right answer). Second, they compare only gross statistical measures of impulse responses (e.g., T30, EDT, D50, C80, TS, G, LF, LFC, IACC) (ISO 3328), and thus it is difficult to determine the causes for simulation errors. Is it because propagation due to edge diffraction is ignored or because angle-dependent reflection functions were not modeled correctly? Without detailed examination of simple impulse responses, these questions are difficult to answer.

Our work is different than most previous validation studies in that we focus only on *simple and controlled acoustical environments*. In contrast to previous simulation systems, we not only use a simple 3D model of the environment, but the environment itself is very simple (e.g., a box with a few configurable panels). In this respect, our work is similar to recent detailed comparisons of geometrical simulations and measurements [30, 32]. These studies show that geometrical room acoustics can lead to satisfying simulations. But Suh and Nelson [30] do not take into account sound diffraction, while Torres et al [32] do not use reverberant environments. In contrast, we present a simulation system and measurement environment which treats the combined effects of specular reflection and diffraction in a reverberant environment. As a result, we are generally able to simulate the sound propagation quite accurately, and we are able to evaluate computed impulse responses with great accuracy. For instance, we can often identify the sequence of surface reflections and edge diffractions that cause a particular peak to appear in an impulse response. Our goal is to obtain a detailed understanding of the limitations of the simulations in these simple environments, and then apply the acquired knowledge to improve our simulations with increased model complexity.

The details of our methods and validation studies appear in the following three sections. First, in Section 3, we explain the algorithms used in our computer simulation system. Then, in Section 4, we describe the construction of the Bell Labs Box and explain how we measure impulse responses. Third, in Section 5, we compare simulations with measurements and discuss our results. While the details of each step are non-trivial, the main message of the paper is that it is only possible to truly understand complicated processes in simple and controlled tests. Only when we have mastered the simple cases (e.g., a box with a diffracting panel) can we consider understanding more complex ones (e.g., a full concert hall).




Figure 3: The four phases of our simulation system.

3 Simulation System

Our sound simulation system is based on the beam tracing method described in [14, 15, 33]. One major advantage of this method over alternative approaches is that it finds all propagation paths with arbitrary combinations of transmission, specular reflection, and diffraction without spatial aliasing. This feature is particularly important for our studies with diffracting panels inside the Bell Labs Box.

The system takes as input: 1) a 3D environment described as a set of polygons with frequency-dependent impedances or bidirectional scattering filters, 2) a speaker described by its location and angular radiation pattern, 3) a microphone described by its location and optional angular directivity pattern. The system outputs a simulated impulse response.

The system executes in four steps, as shown in Figure 3. In the first step, we build a spatial subdivision data structure representing a binary space partition of 3D space into convex polyhedral cells. The purpose of this step is to decompose space into cells whose boundaries are aligned with polygons of the 3D input model and whose adjacencies are stored explicitly in a graph structure in order to enable efficient traversals of 3D space during beam tracing.

In the second step, we trace the convex polyhedral beams representing different propagation sequences through cells of the spatial subdivision. The traversal starts in the cell containing the speaker location with a beam representing the entire cell. Next, it visits adjacent cells iteratively, considering different permutations of transmissions, specular reflections and diffractions due to the faces and edges on the boundary of the “current” cell. As the algorithm traverses a cell boundary into a new cell, a copy of the current convex pyramidal beam is “clipped” to include the region of space passing through the convex polygonal boundary to model transmissions. At each reflecting cell boundary, a copy of the transmission beam is mirrored across the plane supporting the cell boundary to model specular reflections. At each diffracting wedge, a new beam is spawned whose source is the edge and whose extent includes all rays predicted by the Geometrical Theory of Diffraction [20]. The traversal along any sequence terminates when either the length of the shortest path within the beam or the cumulative attenuation exceed some user-specified thresholds. The traversal may also be terminated when the total number of beams traced or the elapsed time exceed other thresholds.

In the third step, for each beam containing the microphone location, we compute the shortest propagation path from the speaker to the microphone along the sequence of transmissions, diffractions, and specular reflections represented by the beam (see Figure 4). The intersections with specularly reflecting faces are uniquely determined by the locations of the speaker, microphone, and the intersections with diffracting edges (diffraction points). The diffrac-

tion points are found by solving a non-linear system of equations expressing equal angle constraints at diffracting edges. Once the diffraction points are found, we construct a piecewise-linear polyline representing the path along which sound travels from source to receiver along the propagation sequence, from which a length-, angle-, and frequency-dependent filter is computed.

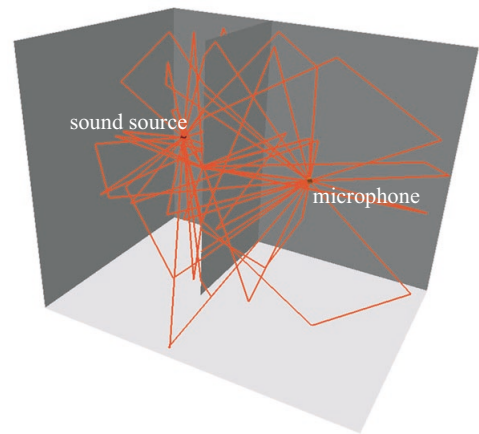


Figure 4: All possible 45 propagation paths combining two specular reflections and one diffraction around the edge of the panel in a model of the Bell Labs Box.

In the final step, for each valid propagation path from the speaker to the microphone, we add its contribution to the simulated impulse response. Our implementation includes source and material filtering effects derived from either measurements (see Section 4) or analytical models. We compute diffraction coefficients using the Uniform Theory of Diffraction [20, 21, 19] (see [33] for details). Atmospheric scattering is also taken into account following the ISO 9013-1 specifications. All calculations are performed in complex Fourier domain at the sampling rate resolution. Our current system uses a sampling rate of 51200 Hz (the sampling rate of our source and material measurements). Thus, for a one second long response we are computing 25600 complex coefficients per path. The simulated transfer function is the sum of the coefficients for *all* paths, and the final impulse response is the inverse Fourier transform of the sum.

After all beams have been traced up to a user-specified termination criterion and the contribution of all propagation paths have been summed, the resulting impulse response is output for comparison to measurements.

4 Sound Measurement Setup

The Bell Labs Box is a small $2.19 \times 3.03 \times 2.42$ meter enclosure (a volume of 16.058 meters^3). It comprises a set of configurable panels, a sound source, a microphone and is constructed using standard residential housing techniques (see Figure 5).




Figure 5: (a) Outside view of the Bell Labs Box. (b) An inside view of the six-sided Bell Labs Box (without diffraction panel).

Sound measurements proceed in two phases. First, in an anechoic chamber, we measure the responses of the sound source, the microphone, and the surface panels independently. Second, we configure the panels into a room and make measurements of the sound propagation from the source to the microphone at different locations within the room that we later compare to simulation results (see Section 5).

Our sound source is a Brüel&Kjaer *artificial mouth* type 4227 speaker [4] (see Figure 6). During the first phase, we measure its directional responses by placing it on a rotator in the Bell Labs anechoic chamber [2], collecting responses at every 5 degrees in the azimuthal plane (this source has a revolution symmetry by design).



Figure 6: Our reference source, the Brüel&Kjaer *artificial mouth*.

Our listening device is a Brüel&Kjaer 4134 1/2-inch microphone, calibrated for free-field recording [4]. Based on the manufacturer's specifications, we assume its frequency response to be flat below 10KHz. We also assume the device to be perfectly omnidirectional below 10KHz. This defines the upper bound on the frequency for which measurement and simulations can be compared. For directional listening devices our system would measure their angle-dependent responses using a setup similar to the one used for sound sources.

Measurements of both the sound source and panel responses are acquired with a *Siglab* measurement system [29] using repeated

chirp stimuli. We use a 51200 Hz sampling rate for all measurements.

The surface panels of the Bell Labs Box are made of 3/4-inch *melamine*, smooth surface facing in, fastened to 3/4-inch FRP plywood panels. All seams are lapped and caulked. Before assembly, we measure their reflection responses using a bidirectional measurement rig in the anechoic chamber (see Figure 7). The rig has six degrees of freedom (three for both source and microphone). Thus, it allows for collecting bidirectional impulse responses of the scattering created by the panel. This setup can be used to measure the acoustic characteristics of any material and derive useful parameters, such as the complex acoustic impedance, which can be considered as an intrinsic material property for hard surfaces. Using the acoustic impedance, complex frequency dependent specular reflection coefficient are derived using the classic model for plane wave reflection which depends only on the impedance and incident angle [23].




Figure 7: Measurement setup for bidirectional material properties in Bell Labs' anechoic chamber. In particular, our six degree of freedom rig allows measurement of the baffle reflection characteristics at different incident directions for source and microphone.

5 Validation Results

In order to evaluate our methods, we ran simulations and measurements for two different geometric panel configurations (see Figure 8). The first is a simple box-shaped enclosure comprising six rectangular panels. The second is the same box-shaped enclosure with a single rectangular panel spanning from floor to ceiling along one half the interior of the box with the speaker and microphone on opposite sides of the panel, as shown in Figure 2. The first configuration is *very simple* yielding only specular reflections, and we use it as a baseline for validation and comparison. The second configuration is a reverberant environment with diffraction and thus incorporates propagation paths combining both edge diffraction and specular reflection. This is a more difficult case for which detail validation results have not previously appeared in the literature.

For both configurations, we measured an impulse response using the Brüel&Kjaer artificial mouth and 1/2-inch microphone connected to the audio outputs and inputs of a *MOTU 828* multi-channel firewire audio interface [27]. The *MOTU* interface was connected to an off-the-shelf laptop running Windows. The source




Figure 9: Comparison of a simulated early impulse response (top) including the first ten orders of specular reflection and a measured response (bottom) in the Bell Labs Box without the baffle. Our simulator computed the contribution of 1524 propagation paths in 738 seconds.




Figure 10: Comparison of a simulated early impulse response (top) including the first two orders of diffraction from the edge of the panel and the first four orders of specular reflection and a measured response (middle) in the Bell Labs Box with the baffle. The simulation computed the contribution of 1358 propagation paths in 631 seconds. The bottom plot shows a simulation including the first eight orders of specular reflection but omitting diffraction (307 paths/153 sec.).




Figure 8: Our two measurement configurations.

signal was a repeated chirp stimuli, and the sampling rate was 48KHz. The output signal used to feed the speaker was also fed back into the interface as a reference. We low-pass filtered the resulting response to get an actual bandwidth of 10KHz. We simulated all possible propagation paths combining up to 10th order specular reflection in the first configuration (empty box), and up to 4th order specular reflection and 2nd order edge diffraction in the second configuration (with diffracting baffle). These simulation parameters were chosen as a reasonable compromise between simulation accuracy and computational expense.

The simulated and measured impulse responses for each configuration are shown in Figures 9 and 10. Note that fine band simulation with measured source and material characteristics allows us to compare the waveform of the simulated and measured early responses. As can be seen, our simulation is able to capture the temporal structure of the impulse response and is also able to capture the effects of the diffraction by the edge of the panel (e.g., first peak in Figure 10), an effect commonly ignored by most acoustic simulation systems. For comparison we also plotted a simulated impulse response without diffraction effects in Figure 10. Omitting diffraction prevents accurate modeling of the early part of the re-

sponse, but has little influence on the late part which is dominated by specular reflection. A significant amount of energy is absent from the early response when reflections alone are simulated. This can introduce errors when evaluating perceptual criteria based on early to late energy ratios (e.g., *clarity*) [3]. We also present in Figure 11 a comparison between spectrograms for the simulated and measured impulse responses of Figures 9 and 10. From these plots, we can conclude that our simulation performs well for the different frequencies. However, differences appear at low frequencies for which geometrical acoustics is no longer a valid model of sound propagation.







Figure 12: In order to track the effect of different propagation sequences on the resulting impulse response, we assign a unique identifier to every edge (cyan) and face (yellow) in our model.




(a) Measured in empty box.



(b) Simulated in empty box.



(c) Measured with baffle.



(d) Simulated with baffle.

Figure 11: Comparison between spectrograms of simulated and measured early responses. The two spectrograms on the left correspond to the (a) measured and (b) simulated impulse responses for the empty box (Figure 9), while those on the right correspond to the (c) measured and (d) simulated responses for the box with a baffle (Figure 10).

Table 1 shows detailed simulation results for the test configuration with the diffracting baffle. Specifically, the columns list the path id, length, time delay, and sequence of scattering events, for the 60 shortest propagation paths (out of 1358 simulated). Every face and edge in the 3D model is assigned a unique identifier (see Figure 12) to allow detailed analysis of each simulated propagation sequence. For instance, looking in the fourth column of the second row, (s 4) denotes a specular reflection off surface 4 and (d 23) denotes a diffraction off edge 23. The ability to relate the propagation delay to the sequence of events along each path is a powerful tool for analyzing the impulse response since it makes it possible to derive correspondences between features in measured and simulated responses. For instance, we labeled some key features of the simulated response in Figure 10 with the corresponding path identifiers in Table 1. The directly diffracted contribution (the shortest possible path) is clearly identified (#1), followed by contributions of the back and side wall occluded by the panel (#2 and #3). Then, a strong specular reflection out of the other side-wall reaches the listener (#5). The strong reflected and diffracted contribution out of the back wall (path #2) is attenuated in our simulation. This might be due to slight inaccuracies in the measured source and microphone positions which results in interference of this path with the diffracted contribution reflected off the ceiling (path #3). This type of detailed analysis makes it possible to explain discrepancies between our simulation results and our measurements, and allows us to improve further our simulation models.

6 Conclusion

In this paper, we have described our initial experiences in building and using the Bell Labs Box, a simple enclosure for validating sound propagation simulations. We find that it is possible to achieve remarkably good matches between simulated and measured impulse responses in a carefully controlled experimental environment. Moreover, since the environment is so simple, we are not only able to validate the simulations at a gross level (as has been done before), but gain insight into the causes for simulation errors from detailed analysis of the differences between simulated and measured responses. We believe such detailed validation is necessary before the results of acoustics simulations can be understood for more complex environments.

7 Future Work

The study described in this paper is just the beginning of a long path toward understanding the validity of computer simulations of acoustical environments.

In the near future, we plan to extend our experiments to more complex environments. Specifically, we intend to consider a series of incremental steps of added complexity (e.g., inserting more surfaces, putting boxes in the enclosure, changing surface materials), measuring and validating the setup after each step. Our general approach is to validate simulations with gradually more complex

Id	Path length (m)	Delay (ms)	Propagation sequence
1	1.6758	4.88	(d 23)
2	2.7205	7.92	(s 4) (d 23)
3	2.7416	7.98	(d 23) (s 7)
4	3.0022	8.74	(s 5) (d 23)
5	3.0154	8.78	(s 8)
6	3.0603	8.91	(d 23) (s 8)
7	3.1522	9.18	(d 23) (s 9)
8	3.1530	9.18	(d 23) (s 10)
9	3.3714	9.82	(s 3) (d 23)
10	3.4798	10.13	(s 4) (s 6) (d 23)
11	3.5356	10.30	(d 23) (s 3) (d 23)
12	3.5356	10.30	(d 23) (s 8) (d 23)
13	3.5645	10.38	(s 4) (s 5) (d 23)
14	3.6860	10.73	(d 23) (s 2)
15	3.7042	10.79	(s 5) (s 6) (d 23)
16	3.7150	10.82	(s 7) (s 8)
17	3.7515	10.92	(d 23) (s 7) (s 8)
18	3.7515	10.92	(s 11) (s 4) (d 23)
19	3.8117	11.10	(s 4) (s 1) (d 23)
20	3.8274	11.15	(d 23) (s 7) (s 10)
21	3.8637	11.25	(s 4) (s 3) (d 23)
22	3.9156	11.40	(d 23) (s 4) (d 23)
23	4.0093	11.68	(s 3) (s 6) (d 23)
24	4.0176	11.70	(s 5) (s 1) (d 23)
25	4.0275	11.73	(s 9) (s 8)
26	4.0612	11.83	(d 23) (s 9) (s 8)
27	4.1051	11.95	(s 4) (d 23) (s 8)
28	4.1315	12.03	(d 23) (s 9) (s 10)
29	4.1483	12.08	(d 23) (s 3) (s 6) (d 23)
30	4.1483	12.08	(d 23) (s 8) (s 7) (d 23)
31	4.1730	12.15	(s 4) (s 5) (s 6) (d 23)
32	4.1956	12.22	(d 23) (s 10) (d 23)
33	4.1956	12.22	(d 23) (s 5) (d 23)
34	4.1978	12.22	(s 4) (d 23) (s 10)
35	4.2772	12.46	(d 23) (s 7) (s 2)
36	4.3005	12.52	(s 3) (s 1) (d 23)
37	4.3338	12.62	(s 11) (s 4) (s 6) (d 23)
38	4.3497	12.67	(s 11) (s 5) (s 4) (d 23)
39	4.3709	12.73	(d 23) (s 8) (s 2)
40	4.3868	12.77	(s 5) (d 23) (s 8)
41	4.4304	12.90	(d 23) (s 3) (s 1) (d 23)
42	4.4308	12.90	(d 23) (s 10) (s 2)
43	4.4313	12.90	(s 4) (s 3) (s 6) (d 23)
44	4.4535	12.97	(s 4) (s 5) (s 1) (d 23)
45	4.4766	13.04	(d 23) (s 4) (s 6) (d 23)
46	4.4795	13.04	(s 5) (d 23) (s 10)
47	4.5513	13.25	(d 23) (s 9) (s 2)
48	4.5804	13.34	(s 4) (d 23) (s 3) (d 23)
49	4.5804	13.34	(s 4) (d 23) (s 8) (d 23)
50	4.5851	13.35	(s 11) (s 4) (s 3) (d 23)
51	4.5872	13.36	(d 23) (s 3) (s 4) (d 23)
52	4.5872	13.36	(d 23) (s 4) (s 3) (d 23)
53	4.6045	13.41	(s 11) (s 4) (s 1) (d 23)
54	4.6433	13.52	(s 4) (d 23) (s 7) (s 8)
55	4.6964	13.68	(s 4) (s 3) (s 1) (d 23)
56	4.7235	13.75	(d 23) (s 5) (s 6) (d 23)
57	4.7254	13.76	(s 4) (d 23) (s 7) (s 10)
58	4.7308	13.78	(s 4) (d 23) (s 2)
59	4.7392	13.80	(d 23) (s 4) (s 1) (d 23)
60	4.7561	13.85	(s 3) (d 23) (s 8)

Table 1: The 60 shortest propagation sequences (out of 1358) including four specular reflections and two diffractions off the edge of the panel in the Bell Labs Box with diffracting baffle. For each path we assign a unique identifier and give the corresponding path length, propagation delay and sequence of events along the path.

environments so that we can understand and quantify the limitations of our simulations in detail. Eventually, we hope to validate simulations of concert halls and other more complex real world environments.

Also, in the near-term, we plan to use our experimental setup to evaluate the accuracy of different models for reflection and diffraction in the Bell Labs Box. For instance, we could compare the Geometrical Theory of Diffraction [20] with the Biot-Medwin-Tolstoy edge diffraction formulation [32], which may lead to more accurate simulations at greater computational expense. Or, we could evaluate more accurate reflection models, such as Thomasson's exact spherical wave reflection model off planar surfaces [31]. The results of these experiments may lead to an understanding of the trade-offs between computational expense and accuracy of different models.

In the longer-term, it will be important to consider psychoacoustical evaluation of acoustic models. We plan to build a system that will use validated acoustic simulations to investigate the psychoacoustic effects of varying acoustic modeling parameters. Our system will allow a user to interactively change acoustics parameters with real-time auralization and visualization feedback. With such a system, it may be possible to address psychoacoustic questions such as: "how many reflections are psychoacoustically important to model?," "which surface reflection model provides a psychoacoustically better approximation?," and "which among conflicting aural and visual cues are dominant in an interactive virtual environment?" We believe that the answers to such questions are of critical importance to future designers of 3D virtual environments.

Acknowledgments

Thomas Funkhouser is partially funded by a National Science Foundation CAREER grant (CCR-0093343) and an Alfred P. Sloan Fellowship. The authors would like to thank the anonymous reviewers for their helpful comments.

References

- [1] Y. Ando. *Concert Hall Acoustics*. Springer-Verlag, 1985.
- [2] Bell Laboratories Anechoic Chamber, <http://www.bell-labs.com/org/1133/research/acoustics/anechoicchamber.html>.
- [3] Leo L. Beranek. *Concert and Opera Halls: How They Sound*. Published for the Acoustical Society of America through the American Institute of Physics, 1996.
- [4] Bruel & Kjaer, USA, <http://www.bkhome.com>.
- [5] J. Blauert. *Spatial Hearing: The Psychophysics of Human Sound Localization*. M.I.T. Press, Cambridge, MA, 1997.
- [6] K.R. Boff, L. Kaufman, and J.P. Thomas, editors. *Handbook of Perception and Human Performance*. Wiley, New York, 1986.
- [7] I. Bork. A comparison of room simulation software - The 2nd round robin on room acoustical computer simulation. *Acustica*, 86(6):943-956, 2000.
- [8] Max Born and Emil Wolf. *Principles of Optics*. 7th edition, Pergamon Press, 1999.
- [9] Bose Corporation, Bose Modeler, Framingham, MA, <http://www.bose.com>.
- [10] CATT-Acoustic, Gothenburg, Sweden, <http://www.netg.se/catt>.
- [11] R.D. Ciskowski and C.A. Brebbia, editors. *Boundary Element Methods in Acoustics*. Elsevier Applied Science, 1991.
- [12] Michael F. Cohen and Donald P. Greenberg. The hemi-cube, a radiosity solution for complex environments. *Computer Graphics (SIGGRAPH 85)*, 19(3):31-40, July 1985.
- [13] N.I. Durlach and A.S. Mavor. Virtual reality scientific and technological challenges. Technical report, National Research Council Report, National Academy Press, 1995.
- [14] T. Funkhouser, I. Carlbom, G. Elko, G. Pingali, M. Sondhi, and J. West. A beam tracing approach to acoustic modeling for interactive virtual environments. *ACM Computer Graphics, SIGGRAPH '98 Proceedings*, pages 21-32, July 1998.
- [15] T. Funkhouser, P. Min, and I. Carlbom. Real-time acoustic modeling for distributed virtual environments. *ACM Computer Graphics, SIGGRAPH '99 Proceedings*, pages 365-374, August 1999.

- [16] C.M. Goral, K.E. Torrance, D.P. Greenberg, and B. Battaile. Modeling the interaction of light between diffuse surfaces. *Computer Graphics (SIGGRAPH 84)*, 18(3), 1984.
- [17] P. Heckbert and P. Hanrahan. Beam tracing polygonal objects. *Computer Graphics (SIGGRAPH 84)*, 18(3):119–127, July 1984.
- [18] J.T. Kajiya. The rendering equation. *ACM Computer Graphics, SIGGRAPH '86 Proceedings*, 20(4), 1986.
- [19] T. Kawai. Sound diffraction by a many sided barrier or pillar. *J. of Sound and Vibration*, 79(2):229–242, 1981.
- [20] J.B. Keller. Geometrical theory of diffraction. *J. of the Optical Society of America*, 52(2):116–130, 1962.
- [21] Robert G. Kouyoumjian and Prabhakar H. Pathak. A uniform geometrical theory of diffraction for an edge in a perfectly conducting surface. *Proc. of IEEE*, 62:1448–1461, November 1974.
- [22] U.R. Krockstadt. Calculating the acoustical room response by the use of a ray tracing technique. *J. Sound and Vibrations*, 8(18), 1968.
- [23] Heinrich Kuttruff. *Room Acoustics (3rd edition)*. Elsevier Applied Science, 1991.
- [24] J. Martin, D. van Maercke, and J.P. Vian. Binaural simulation of concert halls: A new approach for the binaural reverberation process. *J. of the Acoustical Society of America*, 94:3255–3263, December 1993.
- [25] Gary W. Meyer, Holly E. Rushmeier, Michael F. Cohen, Donald P. Greenberg, and Kenneth E. Torrance. An experimental evaluation of computer graphics imagery. *ACM Transactions on Graphics*, 5(1):30–50, January 1986.
- [26] G.R. Moore. *An Approach to the Analysis of Sound in Auditoria*. PhD thesis, Cambridge, UK, 1984.
- [27] Mark Of The Unicorn, USA, <http://www.motu.com>.
- [28] J.M. Naylor. Odeon - another hybrid room acoustical model. *Applied Acoustics*, 38(1):131–143, 1993.
- [29] SigLab, spectral dynamics, USA, <http://www.spectraldynamics.com>.
- [30] J.S. Suh and P.A. Nelson. Measurement of transient responses of rooms and comparison with geometrical acoustic models. *J. of the Acoustical Society of America*, 105(4):2304–2317, April 1999.
- [31] Sven-Ingvar Thomasson. Reflection of waves from a point source by an impedance boundary. *J. of the Acoustical Society of America*, 59(4):780–785, April 1976.
- [32] R. Torres, P. Svensson, and M. Kleiner. Computation of edge diffraction for more accurate room acoustics auralization. *J. of the Acoustical Society of America*, 109:600–610, 2001.
- [33] N. Tsingos, T. Funkhouser, A. Ngan, and I. Carlbom. Modeling acoustics in virtual environments using the uniform theory of diffraction. *ACM Computer Graphics, SIGGRAPH '01 Proceedings*, pages 545–552, August 2001.
- [34] M. Vorlander. International round robin on room acoustical computer simulations. *Proceedings of the 15th International Congress of Acoustics*, June 1995.
- [35] R.B. Welch. *Perceptual Modification: Adapting to Altered Sensory Environments*. Academic Press, New York, 1978.
- [36] Turner Whitted. An improved illumination model for shaded display. *Communications of the ACM*, 23(6):343–349, June 1980.
- [37] E. Zwicker and H. Fastl. *Psychoacoustics: Facts and Models*. Springer, 1999.