

Research Article

Fast Time-Domain Edge-Diffraction Calculations for Interactive Acoustic Simulations

Paul T. Calamia^{1,2} and U. Peter Svensson³

¹Program in Architectural Acoustics, School of Architecture, Rensselaer Polytechnic Institute, Troy, NY 12180, USA

²Department of Computer Science, Princeton University, Princeton, NJ 08544, USA

³Acoustics Research Centre, Department of Electronics and Telecommunications, Norwegian University of Science and Technology, NO-7491 Trondheim, Norway

Received 1 May 2006; Revised 20 September 2006; Accepted 16 October 2006

Recommended by Werner De Bruijn

The inclusion of edge diffraction has long been recognized as an improvement to geometrical-acoustics (GA) modeling techniques, particularly for acoustic simulations of complex environments that are represented as collections of finite-sized planar surfaces. One particular benefit of combining edge diffraction with GA components is that the resulting total sound field is continuous when an acoustic source or receiver crosses a specular-zone or shadow-zone boundary, despite the discontinuity experienced by the associated GA component. In interactive acoustic simulations which include only GA components, such discontinuities may be heard as clicks or other undesirable audible artifacts, and thus diffraction calculations are important for high perceptual quality as well as physical realism. While exact diffraction calculations are difficult to compute at interactive rates, approximate calculations are possible and sufficient for situations in which the ultimate goal is a perceptually plausible simulation rather than a numerically exact one. In this paper, we describe an edge-subdivision strategy that allows for fast time-domain edge-diffraction calculations with relatively low error when compared with results from a more numerically accurate solution. The tradeoff between computation time and accuracy can be controlled with a number of parameters, allowing the user to choose the speed that is necessary and the error that is tolerable for a specific modeling scenario.

Copyright © 2007 P. T. Calamia and U. P. Svensson. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

1. INTRODUCTION

Edge-diffraction impulse responses (IRs) are useful for acoustic simulations involving objects or environments comprising faceted surfaces and have been applied to many problems in acoustics such as loudspeaker radiation [1], noise-barrier analysis [2], and room-acoustics modeling [3]. Diffraction calculations correct for the high-frequency approximation inherent in modeling techniques based on geometrical-acoustics (GA) assumptions, allow for the modeling of sound propagation around occluders and into shadow zones, and provide a smooth, continuous sound-field at specular-zone and shadow-zone boundaries when combined with GA components. All of these factors are important to achieve perceptual realism when auralizing sound fields for virtual-acoustic simulations. When dynamic or interactive simulations are required, continuity of the sound field becomes particularly important. Moving sources or

receivers (i.e., listeners) may cross a zone boundary where the associated GA component (the direct sound or a specular reflection) experiences a discontinuity as it abruptly enters or drops out of the impulse response due to the presence of an occluder or a reflecting surface. Such discontinuities may be heard as clicks or other undesirable audible artifacts if diffraction is not included in the simulated sound field.

Diffraction calculations are generally quite time consuming, a problem that is exacerbated by complex virtual environments with many edges and by interactive simulations which require fast update rates. Existing modeling systems such as those described in [4, 5] address this constraint by using an approach based on the Uniform Theory of Diffraction (UTD) [6], a high-frequency approximation that can be computed quickly. Two assumptions built into the UTD are that the diffracting edge is infinitely long, and that it is far (relative to wavelength) from the source and receiver. An alternative approach which makes no

assumptions about frequency or geometry involves the exact Biot-Tolstoy-Medwin (BTM) expression for diffraction from an infinite rigid wedge [7, 8], which has been derived in a line-integral formulation for finite edges [9]. However, the computational complexity of this method has restricted its use to static scenarios and offline calculations for dynamic simulations.

In this paper, we describe a technique which allows for fast calculations of edge-diffraction impulse responses based on the BTM formulation presented by Svensson et al. in [9]. This formulation is given as an integral along the diffracting edge, suggesting an approach in which the edge can be subdivided into segments for processing. We use a hybrid approach in which each edge is subdivided into two types of segments: sample-aligned segments, each of which contributes to exactly one sample of the diffraction IR; and large evenly sized segments which contribute to multiple IR samples. The former provide a high level of accuracy, but their boundaries are relatively slow to compute and must be updated when the source or receiver is moved. Therefore, we use them only for a small part of the edge which contributes a significant portion of the total diffracted energy to the early part of the IR. The latter segments introduce some error, but their boundaries are independent of the source and receiver positions and can be computed quickly in a preprocessing step for use with the IR tails. The subdivision process, and thus the tradeoff between computation time and accuracy, can be controlled with a number of parameters, allowing the user to choose the speed that is necessary and the error that is tolerable for a specific modeling scenario.

The remainder of this paper is organized as follows. Section 2 discusses related work on diffraction calculations and acoustic modeling. Section 3 contains a brief review of the diffraction formulation in [9] that we use as a basis for our method. Section 4 describes an extension to the edge-subdivision strategy presented in [10] which we use in our system, and Section 5 addresses the various parameters available for adjusting the speed and accuracy of the diffraction calculations. Section 6 presents example calculations along with timing and accuracy data, and Section 7 contains conclusions and suggestions for future work.

2. RELATED WORK

Modeling for interactive acoustic simulations is typically done with one of three basic techniques: the image-source method [11, 12], ray tracing [13], or beam tracing [14].¹ All three are based on geometrical-acoustics assumptions, and thus consider sound propagation only along straight ray-like paths. Such behavior is only correct at asymptotically high frequencies, but GA modeling techniques can provide high levels of accuracy and realism when the dimensions of the

reflecting surfaces are large relative to wavelength. For accurate modeling at low frequencies, with relatively small surfaces, and/or in densely occluded environments, edge diffraction must be taken into account.

While there are many techniques to calculate edge diffraction, two are applied most commonly to acoustic simulations of virtual environments. The first, a frequency-domain method, is the Uniform Theory of Diffraction [6], an extension of the Geometrical Theory of Diffraction [15]. Because the UTD describes diffraction along ray-like paths, it is well suited for integration with GA modeling techniques. UTD-based diffraction can be calculated sufficiently quickly for use in interactive sound-field simulations. However, the UTD is a high-frequency approximation which assumes the diffracting edge has infinite length, and is valid only for source and receiver locations which are far from the edge.

The UTD has been used in two interactive acoustic modeling systems that utilize beam tracing to find the GA components. The first, developed by Funkhouser et al. [4] and Tsingos et al. [16], uses a precomputed beam tree with its root at a fixed source location to identify areas in a 3D model which can be reached by direct, reflected, and/or diffracted sound. As a receiver is moved throughout the modeled environment, the beam tree is used to identify valid propagation paths from the source to the receiver location rapidly, allowing for interactive modeling. For each diffracting edge in a valid path, UTD-based diffraction is calculated based on the shortest path through the edge. For reduced computation time (with a corresponding reduction in accuracy), diffraction calculations can be limited to receivers in shadow zones.

In the second system, developed by Antonacci et al. [5, 17, 18], visibility diagrams for all reflecting surfaces in a 2.5D model (i.e., arbitrarily placed vertical walls with horizontal floors and ceilings) are precomputed using a dual-space representation of the model's geometry. These visibility diagrams allow for rapid construction of beam trees, which in turn allows for interactive modeling with a moving source as well as a moving receiver. UTD diffraction coefficients can be computed for each diffracted path, or can be interpolated from a small set of precomputed values for faster processing.

The second common diffraction-calculation method uses the Biot-Tolstoy-Medwin (BTM) solution, a time-domain formulation for diffraction from a rigid or pressure-release wedge [7, 8, 19]. In particular, the BTM-based expression derived by Svensson et al. in [9] has been used by a number of authors (e.g., see [3, 20, 21]) because it is formulated as a line integral along the diffracting edge and thus is well suited for use with finite edges, and because the BTM solution gives an exact solution for diffraction from a rigid (or pressure release) wedge. Further details of the BTM formulation in [9] are provided in Section 3 as it is the basis for our approximation technique.

The computational complexity of the BTM method has thus far made it of limited use for interactive systems, and this has led to two approximations which were developed specifically to reduce computation time while limiting audible errors in the diffraction. In [22], Lokki et al. calculate diffraction impulse responses using the BTM method

¹ Other acoustic-modeling techniques such as the boundary element method, the finite element method, and various finite-difference schemes attempt to solve the wave equation numerically, but are generally too computationally intensive to allow for interactive calculations and are thus not considered further in this work.

(specifically (2) below), use the FFT to find the diffraction frequency response, and fit a warped infinite-impulse response filter to the smoothed diffraction magnitude response. While this method has been used in dynamic sound rendering with the DIVA system (see [20, 23, 24]), the diffraction calculations must be done offline in a preprocessing stage given the positions of the (moving) source and receiver over time. Interactive simulations are not possible with this approach due to the computation time needed to compute the diffraction IRs and construct the approximation filters. In [25, 26], de Rycker and Torres et al. approximate edge diffraction in a somewhat similar fashion with finite impulse-response (FIR) low-pass filters. However, their method was only tested with static source and receiver pairs, was not integrated into a GA modeling system, and did not provide a way to estimate the frequency-domain diffraction response needed for the filter construction without full impulse-response calculations.

Commercially available acoustic modeling tools such as CATT [27] and Odeon [28] simulate the effects of diffraction on the reflection and scattering from finite surfaces by adjusting the spectra of specular reflections and the fraction of energy that is scattered in nonspecular directions. However, they ignore diffraction into shadow zones, do not calculate explicit diffraction impulse responses, and do not provide interactive simulations. Tsingos and Gascuel [29, 30] interactively simulate occlusion effects due to diffraction from objects between a source and receiver, but also do not calculate diffraction explicitly.

3. BTM EDGE DIFFRACTION

As mentioned above, our diffraction approximations are based on a line-integral formulation of the exact BTM solution as described in [9]. Consider a rigid wedge of finite length, a point source S , and a receiver R whose positions are given with edge-aligned cylindrical coordinates (r_S, θ_S, z_S) and (r_R, θ_R, z_R) , respectively, as shown in Figure 1. The source signal is defined as $q(t) = \rho_0 A(t)/4\pi$, where ρ_0 is the density of air and $A(t)$ is the volume acceleration of the point source. Such a source signal implies that the free-field impulse response of the source is $h(\tau) = \delta(\tau - d/c)/d$, where d is the distance from the source to the receiver and c is the speed of sound. Sound pressure can be calculated through the convolution integral

$$p(t) = \int_0^\infty h(\tau)q(t - \tau)d\tau. \quad (1)$$

The continuous-time edge-diffraction IR at the receiver is given in [9] as an integral over the edge position z ,

$$h(\tau) = -\frac{\nu}{4\pi} \sum_{i=1}^4 \int_{z_1}^{z_2} \delta\left(\tau - \frac{m+l}{c}\right) \frac{\beta_i}{ml} dz, \quad (2)$$

where $\nu = \pi/\theta_W$ is the wedge index, θ_W is the open wedge angle, δ is the Dirac delta function, m and l are the distances from S and R , respectively, to a position on the edge, and c is the speed of sound. The z -coordinate values of the edge




FIGURE 1: Wedge geometry and coordinate system. Locations are specified in cylindrical coordinates where r is the radial distance from the edge, θ is measured from one of the two wedge faces, and the z -axis is aligned with the edge. P_S and P_R are virtual half-planes that contain S and R , respectively, and the edge.

endpoints are used for the integration limits z_1 and z_2 . The functions β_i are

$$\beta_i = \frac{\sin(\nu\varphi_i)}{\cosh(\nu\eta) - \cos(\nu\varphi_i)}, \quad (3)$$

where the angles φ_i are the four combinations of $\pi \pm \theta_S \pm \theta_R$ and the auxiliary function η is

$$\eta = \cosh^{-1} \left\{ \frac{ml + (z - z_S)(z - z_R)}{r_S r_R} \right\}. \quad (4)$$

The shortest path from the source to the receiver through the line that contains the edge goes through the so-called apex point on that line, and this apex point may or may not be contained within the physical edge. If it is, the onset time of the diffraction IR is determined by the path through the apex point. If it is not, the onset time is determined by the shorter of the two paths through the endpoints of the physical edge.

The conversion of (2) into a discrete-time formulation, $h(n)$, can be accomplished by subdividing the edge into segments, and for each segment calculating the IR contribution and distributing it among the appropriate time samples. Numerical integration over each segment is generally straightforward, but the edge-diffraction IR expression is subject to an onset singularity when $\cosh(\nu\eta) = \cos(\nu\varphi_i) = 1$ as seen in (3). This singularity is addressed in [31], and analytical approximations are given for the first sample of the discrete-time IR, $h(n_0)$, which is the only sample affected.

4. EDGE-SUBDIVISION STRATEGIES

Two basic edge-subdivision strategies have been considered previously for the calculation of discrete-time edge-diffraction IRs: subdivision into sample-aligned segments, and

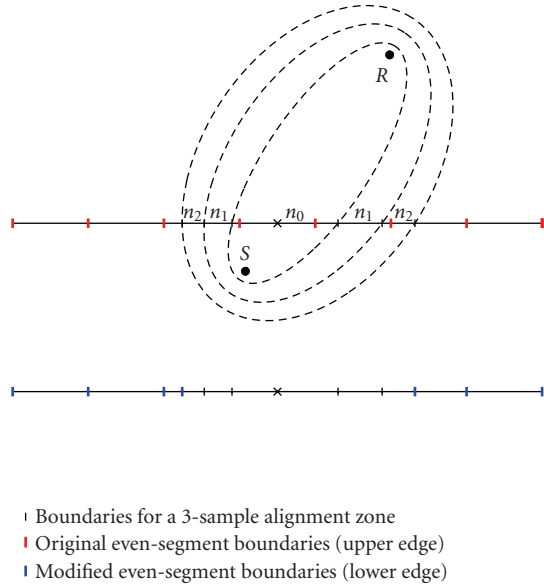


FIGURE 2: Unfolded 2D view of a source, receiver, and segmented edge. The upper edge is marked with the boundaries for a 3-sample alignment zone (samples n_0 , n_1 , and n_2) in black and the original even-segment boundaries in red. The lower edge (S and R not shown) is marked with the modified segment boundaries for the hybrid subdivision scheme in blue: even segments overlapping the alignment zone have been truncated at the edges of the zone, and those completely within the alignment zone have been discarded. The apex point is marked with an “x.”

subdivision into evenly sized segments [9, 10].² A third method, which is a hybrid of these two, was proposed in a simpler form in [10], and the remainder of this paper describes the implementation and the evaluation of a more robust form of that method.

4.1. Subdivision into sample-aligned segments

Sample-aligned segments correspond to portions of an edge which lie between intersections with two confocal ellipsoids (see Figure 2). The foci of the ellipsoids are the source and receiver locations, and the lengths of the axes are determined by the distances $c(n \pm 0.5)/F_S$, where F_S is the sampling frequency and n is the sample index. With such boundaries, each segment contributes to exactly one sample of the discrete-time diffraction IR, which can be written

$$h(n) = -\frac{\gamma}{4\pi} \sum_{i=1}^4 \int_{z_{n,1}}^{z_{n,2}} \frac{\beta_i}{ml} dz. \quad (5)$$

² Medwin et al. [19] and Clay and Kinney [32] also address the conversion of a continuous-time diffraction IR to a discrete-time diffraction IR, although they do so using a form of (2) given as an integral over time so they do not formulate the conversion as an edge-subdivision problem.

Calculation of the integration limits $z_{n,1}$ and $z_{n,2}$ involves finding the roots of a quadratic equation and is described in [31].

Sample-aligned segments are advantageous for many reasons. First, despite the low-pass filtering implied by the area sampling in (5), the spectrum of the discrete-time IR can be made to match that of the exact continuous-time IR up to a chosen frequency by using a sufficiently high sampling rate.³ Second, this method can be used easily with the analytical approximations for sample n_0 in [31] to avoid the onset singularity because the boundaries corresponding to that sample are given explicitly. Finally, the per-segment processing is straightforward: each segment’s contribution is calculated using either numerical integration or the analytical approximation, and the result is added to the corresponding sample of the IR.

Sample-aligned segments unfortunately are not practical for interactive simulations because of the associated computational demands. The segment-boundary calculations are time consuming, the boundaries must be recalculated when the source or receiver is moved, and high sampling frequencies often result in high segment counts.

4.2. Subdivision into evenly sized segments

Evenly sized segments for an edge of length L are generated by choosing a maximum length Δz , and subdividing the edge into k segments of length l where $k = \lceil L/\Delta z \rceil$ and $l = L/k$. The segment-boundary values are easily calculated and are independent of the source and receiver locations; thus they can be calculated once in a simple preprocessing step. However, excessively large values of l or Δz can introduce significant errors in the resulting IR, while small values may result in a prohibitively large number of segments to process. Per-segment processing is somewhat more complicated than with sample-aligned subdivision because each segment may contribute to multiple IR samples. For each segment, the group of corresponding samples must be determined, and the total segment contribution must be calculated and then spread appropriately across these samples. Finally, the boundaries corresponding to sample n_0 are not given explicitly, making it more difficult to avoid the onset singularity.

When using evenly sized segments that contribute to multiple IR samples, the amplitude value A obtained by integrating over the length of a segment must be distributed among the appropriate samples. The path lengths from the source to the receiver through the endpoints of a segment can be used to calculate the span of samples, S_{sp} , to which the segment contributes, and A can be distributed in a number of ways. The simplest approach is to evenly distribute A among the samples, but this leads to a staircase effect in the IR. As described in [10], we apply a correction to the even distribution to achieve a linear approximation of the local slope

³ Clay and Kinney [32] recommend a sampling rate of at least four times the highest frequency of interest for analysis, which suggests using $F_S \geq 80$ kHz for applications in audio and acoustics.

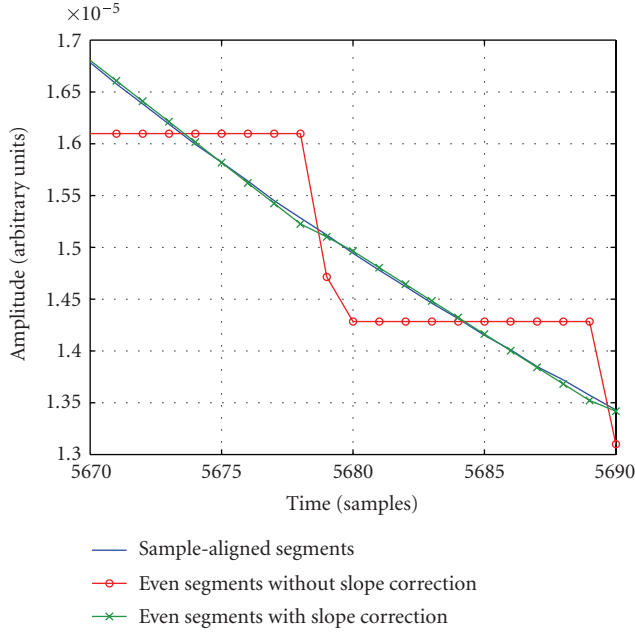


FIGURE 3: Multi-sample distribution for evenly sized segments with and without the slope correction, which assumes the impulse response has a locally linear decay.

of the IR. An example of the multi-sample distribution with and without the slope correction can be seen in Figure 3.

4.3. Hybrid subdivision strategy

A hybrid subdivision strategy can be used to exploit some of the benefits of both sample-aligned and evenly sized segments. With this method, a small number of sample-aligned segments is used to calculate the first N samples of the diffraction impulse response, and evenly sized segments are used to process the remainder of the edge. Any portion of an evenly sized segment that overlaps the alignment zone (i.e., would contribute to any of the first N samples) is discarded. An example of hybrid subdivision with $N = 3$ is shown in Figure 2. If the source or receiver is moved, the boundaries for the first N segments must be recalculated and the alignment-zone overlap tests must be repeated, but this is far less time consuming than recalculating sample-aligned boundaries for the entire edge. This method was introduced in [10], but utilized only a one-sample alignment zone.

5. CALCULATION PARAMETERS

Given the hybrid subdivision method described above, our goal is to minimize the diffraction processing time while limiting the error in the calculations. Three parameters provide control over the accuracy and timing: the number of samples in the alignment zone; the size of the evenly sized segments; and the integration technique used to calculate the contribution of each segment, which can be chosen independently for the alignment zone and the even segments.

5.1. Size of the alignment zone

Because diffraction impulse responses tend to have an impulsive onset followed by a rapidly decaying tail, the high-frequency response is governed by the early samples. The low-frequency response is determined by the total integral over the entire edge, but this value also is strongly dependent on the early part of the IR which has a high amplitude relative to that of the tail. Therefore, accurate computation of the early part of a diffraction IR is critical for an accurate reproduction of its broadband spectral content and thus its perceptual characteristics. Our implementation of the hybrid edge-subdivision scheme allows for an alignment zone of arbitrary size, although as described in Section 6 the use of as few as 4 sample-aligned segments can be sufficient for results with low spectral error.

5.2. Segment size

The size of the even segments is given in terms of the maximum number of IR samples, n_s , to be spanned by any one segment, and converted to a length using

$$\Delta z = \frac{n_s \cdot c}{F_S}, \quad (6)$$

where c is the speed of sound and F_S is the sampling frequency. In practice, the actual sample span of most segments is well below the specified upper bound of n_s . A single value of Δz is used for all edges in a given modeling environment. As Δz increases, computation time decreases due to fewer calls of the integration function, and accuracy decreases because each segment's diffraction contribution must be distributed over a larger span of samples, and the assumption of a locally linear slope over such a span becomes less valid.

5.3. Numerical integration technique

Our implementation provides a choice of three numerical integration techniques: 5-point compound Simpson's rule integration with one step of Romberg extrapolation, standard 3-point Simpson's rule integration, and 1-point midpoint integration [33]. Because the integrand, β_i/ml , includes one hyperbolic and two standard trigonometric functions (see (3)), a reduction of the number of points at which it must be evaluated can lower the total processing time significantly for multi-edge environments, albeit with a corresponding reduction in accuracy. Any of the three techniques can be chosen for the alignment zone and for the evenly sized segments independently. However, the relative importance of the early part of the diffraction IR suggests that combinations in which the integration technique for the alignment zone is equal to or more accurate than that for the evenly sized segments will yield the best results.

6. RESULTS

To evaluate the effect of the parameters described in Section 5 on the computation of diffraction impulse responses, we simulated the diffraction from an array of rigid rectangular




FIGURE 4: Plan view of the panel array used to evaluate the diffraction approximation method. The array was located 5 m above sources S_1 and S_2 and receivers R_1 and R_2 .

panels described in [34], similar to one which might be deployed over the stage in a concert hall or other performance space. The array comprises 35 infinitely thin 1.2 m by 1.2 m panels in a 5-by-7 grid with an inter-panel spacing of 0.5 m. 140 diffracting edges (4 for each panel) were evaluated for each calculation with the array positioned 5 m above two source/receiver pairs. The array and the source and receiver positions are shown in Figure 4. All calculations included first-order diffraction only; neither of the source/receiver pairs engendered a specular reflection from the array, and the direct sound and higher diffraction orders were omitted. Due to the absence of GA components, our testing scenarios are conservative in the sense that they overemphasize the need for accurate diffraction modeling.

All processing was done on a desktop computer with a 3.2 GHz Pentium 4 processor and 2 gigabytes of RAM, and all impulse responses were generated with a sampling rate of 96 kHz. All computation times are averages from 100 trials, and represent the time to compute all 140 diffraction IRs for the total panel-array response. For each of the two source/receiver pairs, the impulse response generated with sample-aligned segments and 5-point integration for all samples was used as the baseline for all speed and accuracy evaluations. Such calculations previously have been shown to agree quite well with measured data [35, 36], so no comparisons to measured data are included here.

For a conservative approximation of the audibility of the errors in the diffraction IRs, we calculated the diffraction magnitude spectra, smoothed them with 1/10-octave filters, and compared them with the smoothed spectrum of the corresponding baseline case. Differences of less than 1 dB between 20 Hz and 20 kHz were assumed to be inaudible and thus acceptable for perceptual accuracy. Even though all diffracting edges were 1.2 m long, the individual impulse responses within the total response from the panel array ranged in size from 16 to 393 samples for the first source/receiver pair, and from 14 to 441 samples for the second source/receiver pair.

Using the hybrid method, we tested 180 combinations of the calculation parameters with each of the two source/

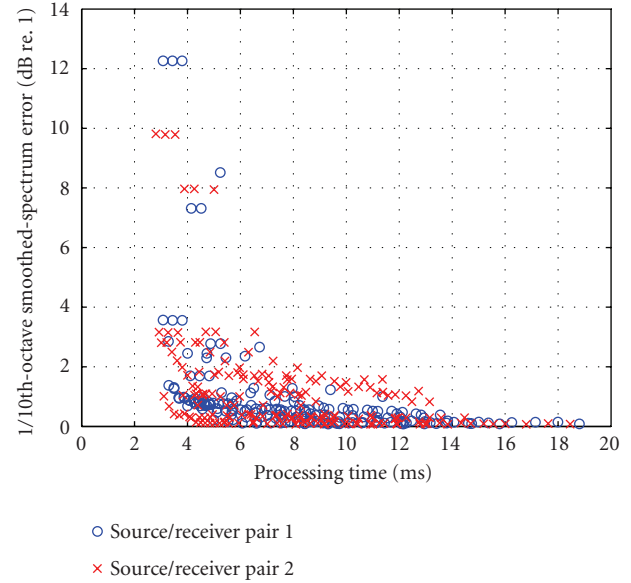


FIGURE 5: Maximum error in the 1/10th-octave smoothed spectra below 20 kHz for various alignment-zone sizes (1 to 10 samples), segment sizes (40, 100, or 300 samples), and integration techniques (1-point, 3-point, or 5-point) using the hybrid method.

receiver pairs. These combinations included: variations in the size of the alignment zone from 1 to 10 samples; three sizes of evenly sized segments, specified as maximum sample spans of 40, 100, and 300 samples; and the three integration techniques used independently on the alignment zone and the evenly sized segments. Only combinations for which the alignment-zone integration technique was equal to, or more accurate than, that for the evenly sized segments were used. For example, given an alignment zone of 5 samples, evenly sized segments limited to a span of no more than 100 samples, and 3-point alignment-zone integration, only 3-point and 1-point integration were tested for the evenly sized segments.

Overall results from the 360 hybrid-subdivision tests are shown in Figure 5, where the maximum error in the 1/10th-octave smoothed spectra (below 20 kHz) for the panel array is plotted against the total processing time. The trend of reduced error with increased processing time is clear, and the effects of the various parameters are generally as expected. For example: all results with a maximum error greater than 4 dB were generated using the largest even-segment size and a single-sample alignment zone; all results with a maximum error less than .09 dB were generated using the smallest segment size and an alignment zone of at least 6 samples; all but one of the results with a processing time less than 4 ms used 1-point integration with 100-sample or 300-sample even segments. Table 1 contains the parameters which resulted in the five fastest processing times with a maximum error of less than 1 dB in the smoothed spectrum for each of the two source/receiver pairs. While there is not a single combination of parameters that yields the “best” result for both source/receiver pairs, the use of a small alignment

TABLE 1: Parameters resulting in the 5 fastest processing times for each S/R pair with maximum error in the smoothed spectrum less than 1 dB. Data for the baseline calculations are also included for comparison as the last entry for each S/R pair.

S/R Pair	Zone Size (Samples)	Zone Integ.	Segment Size (Samples)	Segment Integ.	Proc. Time (ms)	Max. Error (dB)
1	4	1-point	100	1-point	3.67	.97
1	4	1-point	300	1-point	3.69	.94
1	5	1-point	100	1-point	3.86	.98
1	5	1-point	300	1-point	3.88	.96
1	2	3-point	100	1-point	4.01	.69
1	all	5-point	N/A	N/A	171.25	0
2	3	1-point	100	1-point	3.31	.68
2	4	1-point	100	1-point	3.50	.41
2	5	1-point	100	1-point	3.69	.38
2	6	1-point	100	1-point	3.88	.43
2	7	1-point	100	1-point	4.08	.32
2	all	5-point	N/A	N/A	137.51	0

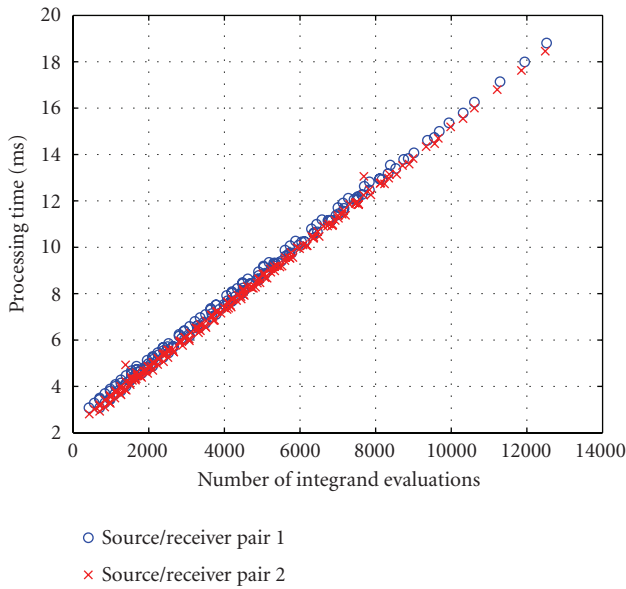


FIGURE 6: Relationship between the processing time and number of evaluations of the diffraction integrand. Processing times are the same as those in Figure 5.

zone ($N \approx 4$ samples) to compute the onset of the diffraction IRs allows for the use of simplified numerical integration and moderately large evenly sized segments and thus rapid calculations with low error. Using the first entry for each source/receiver pair in Table 1, the processing time for S_1 and R_1 was reduced by a factor of 46.6 (from 171 ms to 3.67 ms) and that for S_2 and R_2 by a factor of 41.7 (from 138 ms to 3.31). As can be seen in Figure 6, the processing time for computing the total diffraction impulse response grows linearly with the number of evaluations of the diffraction integrand. This further supports the use of a small align-




FIGURE 7: Impulse responses for source position S_1 and receiver position R_1 . The blue IR (shown in the main figure and the inset) is the baseline calculation using sample-aligned segments and 5-point integration, and the red IR (inset only) is an approximation using hybrid subdivision with the parameters specified in Line 1 of Table 1.

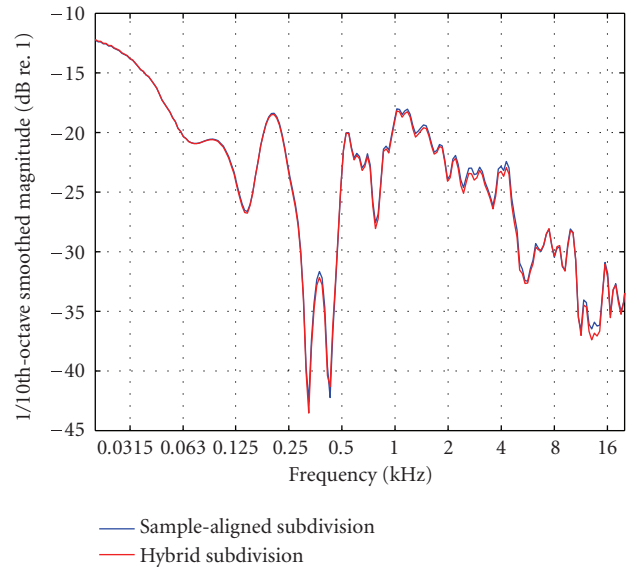


FIGURE 8: 1/10th-octave smoothed magnitude spectra for the impulse responses in Figure 7. See Figure 9 for the difference between the two.

ment zone, moderately large evenly sized segments, and simple integration for rapid calculations.

Results from an example calculation can be seen in Figures 7 and 8. The IRs and corresponding magnitude spectra shown were generated for S_1 and R_1 using the baseline parameter configuration (all sample-aligned subdivision with 5-point integration), and with hybrid subdivision using the

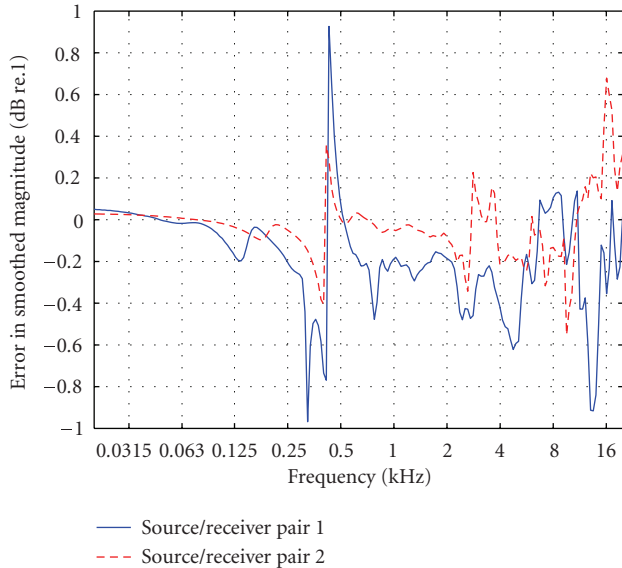


FIGURE 9: Error in the smoothed spectrum for diffraction approximations utilizing the geometry in Figure 4. The solid line is the error for the example shown in Figures 7 and 8 using S_1 and R_1 and the parameters in Line 1 of Table 1. The dashed line is the error for an approximation using S_2 and R_2 and the parameters in Line 7 of Table 1.

following parameters (see Table 1 Line 1): an alignment zone of 4 samples, a maximum sample span for the even segments of 100 samples, and 1-point integration for the entire edge. Figure 7 shows the total impulse response calculated for the panel array, and the inset contains a zoomed-in view of a portion of the IR where the hybrid method's piecewise linear approximation of the IR can be seen. Figure 8 contains the smoothed magnitude spectra of the two IRs. The error (difference between the two spectra) is plotted in Figure 9, as is the error for an example calculation using S_2 and R_2 with the parameters given in Line 7 of Table 1. The maximum error below 20 kHz occurs at approximately 325 Hz for S_1 and R_1 and at 13.05 kHz for S_2 and R_2 .

7. CONCLUSIONS

In this paper, we have presented an edge-subdivision strategy that allows for fast time-domain edge-diffraction calculations with low error. For a given scenario, each edge in a 3D model initially is subdivided into evenly sized segments. As a source and/or receiver is moved in or around the model, the section of each edge which contributes to the first few samples of the edge-diffraction IR, which is dependent on the source and receiver positions, is subdivided into segments such that each one contributes to exactly one sample of the diffraction IR. Even segments overlapping this alignment zone are truncated at the edges of the zone, and those completely within the alignment zone are discarded. The diffraction integral is then evaluated for all remaining segments along the edge. Because the sample-aligned subdivision provides numerically accurate results with a high com-

putational cost, we generally restrict its use to the early portion of the IR which contains a significant percentage of the total diffracted energy and thus must be calculated accurately for a perceptually convincing simulation. However, a user is free to choose the size of the alignment zone, as well as the size of the evenly sized segments and the integration technique(s) used to evaluate the two types of segments to optimize the speed-accuracy tradeoff for each modeling scenario. Tests on an array of rigid panels indicate that fast, low-error results can be obtained with an alignment zone of approximately 4 samples, even segments limited to a 100-sample span (with a sampling frequency of 96 kHz), and numerically simple midpoint integration for all segments. Our method can be combined easily with a geometrical-acoustics modeling approach such as image-source modeling or beam tracing to provide rapid calculations of a smooth continuous sound field for interactive acoustic simulations.

This work suggests several avenues for future research. First, we would like to study the audibility of the errors introduced by the subdivision strategy. Our error criterion of 1 dB in the 1/10th-octave smoothed spectra is quite conservative and can likely be relaxed significantly when the diffraction calculations are combined with GA components which typically dominate the sound field. However, listening tests could provide more definitive results, and possibly could provide guidelines for choosing subdivision parameter values appropriate for different applications. The listening tests also could include comparisons to other fast diffraction approximations, for example, those based on the UTD. While our method should provide higher accuracy, the perceptual benefits of that accuracy have yet to be established.

Second, the extension of this method to higher orders of diffraction and nonrigid surfaces should also be considered. In regard to the former, there is some evidence to suggest that higher orders of diffraction can be neglected when auralizing simulated impulse responses [3]. However, this has not been tested in interactive or dynamic simulations in which discontinuities in the geometrical-acoustics field may result in audible artifacts. In regard to nonrigid surfaces, solutions for diffraction from a density-contrast wedge have been developed [37, 38], but there currently is no known general solution for diffraction from a wedge of arbitrary impedance. Many interior and exterior building materials (e.g., concrete, brick, and thick plaster) which are likely to be simulated in virtual environments are acoustically "hard," thus the assumption of rigidity may not be a significant limitation. Nevertheless, a more general edge-diffraction expression applicable to all materials would allow for more accurate results, and edge-subdivision could be applied to it for rapid calculations.

Third, to reduce computation time further, we could consider a number of ways to omit edges during processing. For example, we could use an approach similar to the one in [16] in which diffraction calculations can be limited to edges for which the receiver is in the shadow zone. However, it might be helpful to extend the calculation range to include edges for which the receiver is within a small angular distance of a zone boundary. This would ensure sound-field continuity across all zone boundaries, and would capture the

diffraction IRs where their amplitudes are maximal. Another approach could involve omitting edges for which the apex point is not included within the edge since the amplitude of the diffraction IR for such cases is small. Preliminary results for these forms of “diffraction culling” are presented in [39].

Finally, this subdivision scheme may aid in two additional aspects of virtual acoustic simulations that include diffraction: visibility calculations and directional auralization. In regard to visibility, simulations of densely occluded environments require additional calculations to determine the portion(s) of each edge to which the source and receiver both have clear lines of sight. Approximate visibility could be computed by simply considering the visibility of the midpoint of each segment from the source and receiver. In regard to auralization of diffraction, Torres et al. [3] have suggested convolving each diffraction IR with the head-related impulse response specific to the direction of the least-time diffracted path. In some cases, however, an edge (from apex point to end point) may subtend an angle large enough such that a single direction of arrival for the entire IR is not sufficient for a perceptually accurate simulation. In such cases, the diffraction IR could be auralized using paths through the midpoints of the segments (or a subset of segments) to allow the direction of arrival to change over time.

ACKNOWLEDGMENTS

The authors would like to thank the Norway-America Association, the Acoustics Research Centre project from the Research Council of Norway, and the National Science Foundation (Grant CCR-0093343) for partially funding this work, as well as the anonymous reviewers and the Associate Editor, Werner de Bruijn, for their helpful comments and suggestions.

REFERENCES

- [1] J. Vanderkooy, “A simple theory of cabinet edge diffraction,” *Journal of the Audio Engineering Society*, vol. 39, no. 12, pp. 923–933, 1991.
- [2] P. Menounou and J. H. You, “Experimental study of the diffracted sound field around jagged edge noise barriers,” *The Journal of the Acoustical Society of America*, vol. 116, no. 5, pp. 2843–2854, 2004.
- [3] R. R. Torres, U. P. Svensson, and M. Kleiner, “Computation of edge diffraction for more accurate room acoustics auralization,” *The Journal of the Acoustical Society of America*, vol. 109, no. 2, pp. 600–610, 2001.
- [4] T. Funkhouser, N. Tsingos, I. Carlbom, et al., “A beam tracing method for interactive architectural acoustics,” *The Journal of the Acoustical Society of America*, vol. 115, no. 2, pp. 739–756, 2004.
- [5] F. Antonacci, M. Foco, A. Sarti, and S. Tubaro, “Fast modeling of acoustic reflections and diffraction in complex environments using visibility diagrams,” in *Proceedings of 12th European Signal Processing Conference (EUSIPCO '04)*, pp. 1773–1776, Vienna, Austria, September 2004.
- [6] R. G. Kouyoumjian and P. H. Pathak, “A uniform geometrical theory of diffraction for an edge in a perfectly conducting surface,” *Proceedings of the IEEE*, vol. 62, pp. 1448–1461, 1974.
- [7] M. A. Biot and I. Tolstoy, “Formulation of wave propagation in infinite media by normal coordinates with an application to diffraction,” *The Journal of the Acoustical Society of America*, vol. 29, no. 3, pp. 381–391, 1957.
- [8] H. Medwin, “Shadowing by finite noise barriers,” *The Journal of the Acoustical Society of America*, vol. 69, no. 4, pp. 1060–1064, 1981.
- [9] U. P. Svensson, R. I. Fred, and J. Vanderkooy, “An analytic secondary source model of edge diffraction impulse responses,” *The Journal of the Acoustical Society of America*, vol. 106, no. 5, pp. 2331–2344, 1999.
- [10] P. T. Calamia and U. P. Svensson, “Edge subdivision for fast diffraction calculations,” in *Proceedings of IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*, pp. 187–190, New Paltz, NY, USA, October 2005.
- [11] J. B. Allen and D. A. Berkley, “Image method for efficiently simulating small-room acoustics,” *The Journal of the Acoustical Society of America*, vol. 65, no. 4, pp. 943–950, 1979.
- [12] J. Borish, “Extension of the image model to arbitrary polyhedra,” *The Journal of the Acoustical Society of America*, vol. 75, no. 6, pp. 1827–1836, 1984.
- [13] A. Krokstad, S. Strøm, and S. Sørsdal, “Calculating the acoustical room response by the use of a ray tracing technique,” *Journal of Sound and Vibration*, vol. 8, no. 1, pp. 118–125, 1968.
- [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,” in *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '98)*, pp. 21–32, Orlando, Fla, USA, July 1998.
- [15] J. B. Keller, “Geometrical theory of diffraction,” *Journal of the Optical Society of America*, vol. 52, no. 2, pp. 116–130, 1962.
- [16] N. Tsingos, T. Funkhouser, A. Ngan, and I. Carlbom, “Modeling acoustics in virtual environments using the uniform theory of diffraction,” in *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '01)*, pp. 545–552, Los Angeles, Calif, USA, August 2001.
- [17] F. Antonacci, M. Foco, A. Sarti, and S. Tubaro, “Accurate and fast audio-realistic rendering of sounds in virtual environments,” in *Proceedings of 6th IEEE Workshop on Multimedia Signal Processing*, pp. 271–274, Siena, Italy, September–October 2004.
- [18] F. Antonacci, M. Foco, A. Sarti, and S. Tubaro, “Real time modeling of acoustic propagation in complex environments,” in *Proceedings of 7th International Conference on Digital Audio Effects (DAFx '04)*, pp. 274–279, Naples, Italy, October 2004.
- [19] H. Medwin, E. Childs, and G. M. Jebsen, “Impulse studies of double diffraction: a discrete Huygens interpretation,” *The Journal of the Acoustical Society of America*, vol. 72, no. 3, pp. 1005–1013, 1982.
- [20] V. Pulkki, T. Lokki, and L. Savioja, “Implementation and visualization of edge diffraction with image-source method,” in *Proceedings of 112th Audio Engineering Society Convention*, Munich, Germany, May 2002.
- [21] P. T. Calamia, U. P. Svensson, and T. Funkhouser, “Integration of edge-diffraction calculations and geometrical-acoustics modeling,” in *Proceedings of Forum Acusticum*, pp. 2499–2504, Budapest, Hungary, August 2005.
- [22] T. Lokki, U. P. Svensson, and L. Savioja, “An efficient auralization of edge diffraction,” in *Proceedings of the Audio Engineering Society 21st International Conference on Architectural*

Acoustics and Sound Reinforcement, pp. 166–172, St. Petersburg, Russia, June 2002.

- [23] L. Savioja, J. Huopaniemi, T. Lokki, and R. Väänänen, “Creating interactive virtual acoustic environments,” *Journal of the Audio Engineering Society*, vol. 47, no. 9, pp. 675–705, 1999.
- [24] L. Savioja, J. Huopaniemi, and T. Lokki, “Auralization applying the parametric room acoustic modeling technique—the DIVA auralization system,” in *Proceedings of the 8th International Conference on Auditory Display*, Kyoto, Japan, July 2002.
- [25] N. de Rycker, “Theoretical and numerical study of sound diffraction—application to room acoustics auralization,” *Rapport de Stage D’Option Scientifique*, École Polytechnique, Paris, France, 2002.
- [26] R. Torres, N. de Rycker, and M. Kleiner, “Edge diffraction and surface scattering in concert halls: physical and perceptual aspects,” *Journal of Temporal Design in Architecture and the Environment*, vol. 4, pp. 52–58, 2004.
- [27] B.-I. Dalenbäck, *CATT-Acoustic v8 Manual*, <http://www.catt.se/>.
- [28] C. L. Christensen, *ODEON Room Acoustics Program ver. 8 Manual*, <http://www.odeon.dk>.
- [29] N. Tsingos and J.-D. Gascuel, “Soundtracks for computer animation: sound rendering in dynamic environments with occlusions,” in *Proceedings of the Conference on Graphics Interface*, pp. 9–16, Kelowna, British Columbia, Canada, May 1997.
- [30] N. Tsingos and J.-D. Gascuel, “Fast rendering of sound occlusion and diffraction effects for virtual acoustic environments,” in *Proceedings of the Audio Engineering Society 104th Convention*, Amsterdam, The Netherlands, May 1998, preprint no. 4699.
- [31] U. P. Svensson and P. T. Calamia, “Edge-diffraction impulse responses near specular-zone and shadow-zone boundaries,” *Acta Acustica united with Acustica*, vol. 92, no. 4, pp. 501–512, 2006.
- [32] C. S. Clay and W. A. Kinney, “Numerical computations of time-domain diffractions from wedges and reflections from facets,” *The Journal of the Acoustical Society of America*, vol. 83, no. 6, pp. 2126–2133, 1988.
- [33] P. J. Davis and P. Rabinowitz, *Methods of Numerical Integration*, Academic Press, New York, NY, USA, 2nd edition, 1984.
- [34] R. Torres, *Studies of edge diffraction and scattering: applications to room acoustics and auralization*, Ph.D. thesis, Chalmers University of Technology, Göteborg, Sweden, 2000.
- [35] T. Lokki and V. Pulkki, “Measurement and theoretical validation of diffraction from a single edge,” in *Proceedings of the 18th International Congress on Acoustics (ICA '04)*, vol. 2, pp. 929–932, Kyoto, Japan, April 2004.
- [36] A. Løvstad and U. P. Svensson, “Diffracted sound field from an orchestra pit,” *Acoustical Science and Technology*, vol. 26, no. 2, pp. 237–239, 2005.
- [37] A. M. J. Davis and R. W. Scharstein, “The complete extension of the Biot-Tolstoy solution to the density contrast wedge with sample calculations,” *The Journal of the Acoustical Society of America*, vol. 101, no. 4, pp. 1821–1835, 1997.
- [38] J. C. Novarini and R. S. Keiffer, “Impulse response of a density contrast wedge: practical implementation and some aspects of its diffracted component,” *Applied Acoustics*, vol. 58, no. 2, pp. 195–210, 1999.
- [39] U. P. Svensson and P. T. Calamia, “The use of edge diffraction in computational room acoustics,” *The Journal of the Acoustical Society of America*, vol. 120, p. 2998, 2006, (A).

Paul T. Calamia is an Assistant Professor in the Graduate Program in Architectural Acoustics in the School of Architecture at Rensselaer Polytechnic Institute in Troy, NY. He earned a B.S. degree in mathematics from Duke University (1992), and an M.S. degree in electrical engineering in the Engineering Acoustics Program at the University of Texas at Austin (1998). He is currently in the process of completing a Ph.D. degree in computer science at Princeton University, and he spent one year as a guest Ph.D. student with the Acoustics Group at the Norwegian University of Science and Technology in Trondheim. Between degrees, he spent two years working as an Acoustical Engineer with Wyle Laboratories in Arlington, VA, and four years working as an Acoustics Consultant with Kirkegaard Associates in Chicago, IL. He is an Active Member of the Acoustical Society of America and its Technical Committee on Architectural Acoustics. His research interests include computational room acoustics, interactive sound-field simulations, edge-diffraction modeling, and spatial audio rendering.



U. Peter Svensson received his M.S. degree (Eng. Phys.) in 1987 and his Ph.D. degree in 1994, both from Chalmers University of Technology, Gothenburg, Sweden. He has held postdoctoral positions at Chalmers University, University of Waterloo, Ontario, Canada, and Kobe University, Japan. Since 1999, he has been a Professor in electroacoustics at NTNU in Trondheim. His research interests are auralization and sound reproduction techniques, time-domain models of sound radiation and scattering, measurement techniques, and audio over packet networks. He has published 18 journal papers and more than 60 conference papers. He has been a Board Member of the Acoustical Societies of Sweden and Norway, including a period as chairman for the latter. He is Associate Editor in the field of computational acoustics for *Acta Acustica united with Acustica*. In 2001, he received an award for the best contribution by an author 35 years or younger to the *Journal of Audio Engineering Society* for the paper “Errors in MLS measurements caused by time variance in acoustic systems” together with Dr. Johan L. Nielsen. He is a Member of the Acoustical Societies of America and Japan and the Audio Engineering Society.

