Simple conformal methods for FDTD modeling of pressure-release surfaces

Abbreviated title: Conformal Pressure-Release FDTD Methods

PACS numbers: 43.20Fn, 43.30Hw, 43.30Gv, 43.20Px

John B. Schneider, email: schneidj@eecs.wsu.edu
Christopher L. Wagner, email: clwagner@eecs.wsu.edu
School of Electrical Engineering and Computer Science
Washington State University
P.O. Box 642752
Pullman, WA 99164-2752
(509) 335-4655, FAX (509) 335-3818

Robert J. Kruhlak, email: rkruhlak@wsu.edu
Nonlinear Optics Lab., Department of Physics
Washington State University
P.O. Box 642814
Pullman, WA 99164-2814

Received:

Abstract

The finite-difference time domain (FDTD) method provides a simple and accurate means of simulating a wide range of acoustic wave propagation problems. Unfortunately, the method has a voracious appetite for computational resources. For example, to accurately model scattering from a continuously varying pressure-release boundary, an FDTD grid is typically required that has a much finer discretization than is necessary to model propagation in a homogeneous space. Such a fine discretization can become prohibitive when considering large-scale problems. We present two simple conformal techniques for acoustic FDTD simulations of problems involving pressure-release boundaries. These techniques, which rely upon splitting velocity cells adjacent to the pressure-release boundary, significantly improve the accuracy of the results over those of the standard "staircase" representation of the boundary. These methods permit the use of a coarser FDTD grid than would otherwise be practical and yet add negligible computational cost. The improved accuracy of these split-cell techniques is shown for both a spherical scatterer and a spherical resonator.

Introduction

The finite-difference time-domain (FDTD) method was introduced by Yee in 1966 to study electromagnetic wave propagation [1]. The method employs a discretization of time and space to express the coupled first-order Maxwell's equations as difference equations. From these difference equations, the future fields are expressed in terms of known past fields which yields a leap-frog technique for the advancement of the fields. A similar method was later presented by Madariaga [2] and Virieux [3, 4] for elastic wave propagation where the relevant fields are stress and velocity. A three-dimensional version of the stress-velocity FDTD method was presented by, among others, Randall [5].

The FDTD simulation of acoustic propagation can be considered a special case of the elastic stress-velocity method and subsequent references to the FDTD method imply the acoustic version where the fields being modeled are pressure and velocity. The FDTD method has been used to solve a wide range of problems including the analysis of room acoustics [6, 7] and the determination of acoustic scattering from a rough ocean surface (i.e., a pressure-release surface) [8, 9]. The implementation details for the method can be found in several papers (e.g., [7, 10, 11, 12]).

Because the FDTD method is a full-wave solution that directly employs the governing differential equations, the solution of three-dimensional problems is often computationally costly. The accuracy of the FDTD method is, in large part, tied to the fineness of the computational grid. When a fine grid must be used and the computation space is large, the computational resources needed to obtain the solution may be prohibitive. Some problems, such as those with planar boundaries which are aligned with the grid, can be solved with relatively coarse grid. However, the simulation of more general problems, such as those with continuously varying surfaces, may suffer intolerable numerical errors when modeled with a coarse grid. The goal of the conformal methods presented here is to provide an accurate representation of pressure-release surfaces even when the grid is relatively coarse and the surface is not aligned with the grid. Thus, these methods serve to reduce the error introduced by the "staircase" representation of material boundaries.

The methods introduced here are designed to improve the geometric fidelity of simulations involving pressure-release surfaces. These methods do not address the numerical dispersion errors which are inherent in the FDTD grid and which increase as the grid becomes more coarse. However, we note that there are ways to mitigate dispersion errors and, consequently, in many simulations the staircasing errors (and not the dispersion errors) dictate the acceptable coarseness of the grid. For example, in [13] the FDTD dispersion relation was used to correct for dispersion errors. Another example is provided by the work of Hastings *et al.* [9] where the incident field was introduced, and the scattered field was recorded, near the scatterer. In this way, the fields do not propagate far within the grid and, therefore, do not accumulate significant dipsersion error.

In the standard FDTD method, the velocity and pressure nodes are at fixed locations in a Cartesian grid so that a continuously varying boundary is approximated by one that is staircased. Furthermore, even flat planar surfaces must be approximated as a staircased surface when the plane is not aligned with the grid. With a sufficiently fine grid, the errors associated with the staircase approximation are small. However, to keep the computational cost as low as possible, the grid should be made as coarse as possible.

In an attempt to satisfy these competing factors, different approaches have been proposed. For example, Botteldooren presented a quasi-Cartesian scheme that can be used to distort the grid so that it conforms to a boundary [14]. Unfortunately, such a scheme removes much of the simplicity and elegance of the Cartesian-grid FDTD method and the generation of quasi-Cartesian grids for arbitrary scatterers is a significant computational task. An alternative approach to reducing the errors associated with the staircase representation is to maintain the Cartesian grid away from material boundaries and to distort the grid only in the vicinity of the boundaries. Such an approach was put forward in the electromagnetics literature by, among others, Jurgens *et al.* [15] and it has been used to model scattering from two-dimensional pressure-release surfaces [9]. Unfortunately, for three-dimensional problems, there is no clear translation of this method to the acoustic case.

We present two simple methods that alleviate most of the errors associated with the staircase representation of pressure-release surfaces. The techniques, which have their antecedents in the work of Mezzanotte *et al.* [16] and Dey and Mittra [17], are realized by splitting cells associated with velocity nodes. In the simplest version, the mass of a cell is reduced by a factor of two if the pressure-release boundary is within one quarter of the spatial step distance to either side of a velocity node. This simple modification significantly improves the accuracy as compared to the standard staircase solution and, furthermore, has associated with it virtually no additional computational cost. The simulation can still be run at the Courant limit (i.e., using the greatest possible time step as dictated by the size of the Cartesian grid away from the boundary).

The simplicity of these techniques is, in part, a consequence of only considering pressure-release surfaces and not boundaries between arbitrary materials. Nevertheless, it is envisioned that these techniques will make possible, for example, the investigation of scattering from realistic ocean surfaces (i.e., a three-dimensional problem with a two-dimensional ocean surface) via the FDTD method. Additionally, a similar sort of cell-splitting (or mass averaging) technique is possible for more general problems as was shown in [18] where the scattering from fluid-fluid interfaces was investigated. However, the more general case is not considered here.

The following section presents the details of the conformal methods for acoustic pressure-release boundaries. Results are then presented in Sec. 2 for a spherical pressure-release scatterer and for a spherical resonator.

1 Conformal Techniques

The governing acoustic equations in three dimensions are

$$\frac{\partial p}{\partial t} = -\rho c^2 \left(\frac{\partial v_x}{\partial x} + \frac{\partial v_y}{\partial y} + \frac{\partial v_z}{\partial z} \right), \tag{1}$$

$$\frac{\partial v_x}{\partial t} = -\frac{1}{\rho}\frac{\partial p}{\partial x}, \tag{2}$$

$$\frac{\partial v_y}{\partial t} = -\frac{1}{\rho}\frac{\partial p}{\partial y}, \tag{3}$$

$$\frac{\partial v_z}{\partial t} = -\frac{1}{\rho}\frac{\partial p}{\partial z}, \tag{4}$$

where p is pressure, $\vec{v} = v_x \hat{\mathbf{a}}_x + v_y \hat{\mathbf{a}}_y + v_z \hat{\mathbf{a}}_z$ is velocity, ρ is density, and c is the speed of propagation. The update equations for the FDTD method are obtained by replacing the derivatives in Eqs. (1)-(4) by finite differences and solving for unknown future fields in terms of known past fields. To obtain a fully explicit scheme, the points at which the fields are evaluated must be offset spatially and temporally. Second-order accurate central differences can be used to approximate all the derivatives if the fields are discretized so that they are defined at the following evaluation points:

$$p(x, y, z, t) = p(i\Delta x, j\Delta y, k\Delta z, n\Delta t) = p^n(i, j, k), \tag{5}$$

$$v_x(x, y, z, t) = v_x\left([i + 1/2]\Delta x, j\Delta y, k\Delta z, [n + 1/2]\Delta t\right) = v_x^{n+1/2}(i, j, k), \tag{6}$$

$$v_y(x,y,z,t) = v_y\left(i\Delta x, [j+1/2]\Delta y, k\Delta z, [n+1/2]\Delta t\right) = v_y^{n+1/2}(i,j,k), \tag{7}$$

$$v_z(x,y,z,t) = v_z\left(i\Delta x, j\Delta y, [k+1/2]\Delta z, [n+1/2]\Delta t\right) = v_z^{n+1/2}(i,j,k), \tag{8}$$

where Δx, Δy, and Δz are the spatial step sizes in the x, y, and z directions, respectively, and Δt is the temporal step size. The precise interpretation of a field value in an FDTD grid is an interesting topic in itself (e.g., are the fields samples of a continuous function or do they represent average values over some region of space?), and the interested reader is referred to [19] for further discussions. For purposes of this paper, a field is assumed to be defined at a point (i.e., a particular node in the grid) and the value at the point represents the average value found over the cell associated with that field component. The arrangement of nodes associated with a given set of spatial indices is shown in Fig. 1. We assume a uniform grid so that $\Delta x = \Delta y = \Delta z = \delta$, but the development of the conformal methods do *not* require this assumption. On the right-hand side of Eqs. (5)-(8), the fields are specified by their spatial and temporal indices, i.e., the arguments and the superscripts, respectively. For the velocity components, the spatial offsets are implied by the field component (i.e., the x component of velocity $v_x^{n+1/2}(i,j,k)$ is not collocated with either the y component $v_y^{n+1/2}(i,j,k)$ or z component $v_z^{n+1/2}(i,j,k)$). The temporal offset between the fields is explicitly retained in the temporal index. Replacing the derivatives in Eqs. (1)-(4) with finite differences and using the discretization of Eqs. (5)-(8) yields the following update equations:

$$\begin{aligned} p^n(i,j,k) = p^{n-1}(i,j,k) - \rho c \frac{c\Delta t}{\delta} \Big[& v_x^{n-1/2}(i,j,k) - v_x^{n-1/2}(i-1,j,k) + \\ & v_y^{n-1/2}(i,j,k) - v_y^{n-1/2}(i,j-1,k) + \\ & v_z^{n-1/2}(i,j,k) - v_z^{n-1/2}(i,j,k-1) \Big], \end{aligned} \tag{9}$$

$$v_x^{n+1/2}(i,j,k) = v_x^{n-1/2}(i,j,k) - \frac{1}{\rho c}\frac{c\Delta t}{\delta}\left[p^n(i+1,j,k) - p^n(i,j,k)\right], \tag{10}$$

$$v_y^{n+1/2}(i,j,k) = v_y^{n-1/2}(i,j,k) - \frac{1}{\rho c}\frac{c\Delta t}{\delta}\left[p^n(i,j+1,k) - p^n(i,j,k)\right], \tag{11}$$

$$v_z^{n+1/2}(i,j,k) = v_x^{n-1/2}(i,j,k) - \frac{1}{\rho c}\frac{c\Delta t}{\delta}\left[p^n(i,j,k+1) - p^n(i,j,k)\right]. \tag{12}$$

These equations are used in a leap-frog fashion to obtain the unknown future fields in terms of the known past fields. This explicit scheme is not unconditionally stable. To obtain a stable solution in three dimensions, the Courant number, i.e., $c\Delta t/\delta$, must be less than or equal to $1/\sqrt{3}$. (For more information on the Courant number and stability of the FDTD method see [10].)

The simplest staircase implementation of a pressure-release region is realized by setting all pressure nodes that lie within the pressure-release region to zero. Pressure nodes outside the pressure-release region are updated using Eq. (9) and all velocity nodes, whether they fall inside or outside the pressure-release region, are updated as dictated by Eqs. (10)-(12).

Figure 2 shows a cross-section of the cell (or cube) associated with the velocity node $v_x(i,j,k)$ (the full cell extends $\Delta y/2$ into and out of the page about the velocity node). Assume that a pressure-release boundary is normal to the x direction at a distance $L\Delta x$ from the node $p(i,j,k)$, where $0 < L \leq 1$ and that the pressure-release region is to the right of this boundary. The simplest staircase implementation of this boundary is realized by setting $p(i+1,j,k)$ to zero while $p(i,j,k)$ and $v_x(i,j,k)$ are updated in the usual way for all values of L. However, this results in an approximation of the actual boundary that suffers unnecessarily large distortions for L's such that $0 < L < 1/2$. Instead, the node $p(i,j,k)$, which is the closest pressure node to the boundary when $0 < L < 1/2$, should be set to zero. Thus an alternative staircase representation of the boundary can be obtained by changing the criterion used to set a pressure node to zero. Instead of setting to zero those pressure nodes that lie within the pressure-release region (i.e., beyond the pressure-release boundary), pressure nodes should be set to zero if any of their neighboring velocity nodes lie within the pressure-release region. Thus, $p(i,j,k)$ is set to zero if any of the following velocities are within the pressure-release region: $v_x(i-1,j,k)$, $v_x(i,j,k)$, $v_y(i,j-1,k)$, $v_y(i,j,k)$, $v_z(i,j,k-1)$, or $v_z(i,j,k)$. We label the staircasing technique that is based solely on the location of the pressure nodes as the Stair-p method. The staircasing technique which tests the location of the neighboring velocity nodes is labeled Stair-v.

In both staircasing methods all the velocity nodes, independent of their location, are updated in the usual way. The velocity nodes "see" the pressure-release boundary only through the pressure nodes that have been set to zero. The velocity update equations are obtained from Newton's equation of motion. Ignoring for the moment any pressure-release boundaries, the mass of the cell associated with velocity node $v_x(i,j,k)$ is $\rho\Delta x\Delta y\Delta z$, the force on the cell in the positive x direction is $p(i,j,k)\Delta y\Delta z$, the force in the opposite direction is $p(i+1,j,k)\Delta y\Delta z$, and the acceleration in the x direction is $\partial v_x(i,j,k)/\partial t$. Setting the mass times the acceleration equal to the sum of the forces yields

$$\rho\Delta x\Delta y\Delta z\frac{\partial v_x(i,j,k)}{\partial t} = -\Delta y\Delta z\left(p(i+1,j,k) - p(i,j,k)\right). \tag{13}$$

Expressing the temporal derivative as a central difference and solving for $v_x^{n+1/2}(i,j,k)$ yields Eq. (10). When the pressure-release boundary bisects a cell, i.e., when the boundary is located at $L = 1/2$, the mass of the cell is reduced by one half. Using this mass and again expressing the temporal derivative as a central difference, Eq. (13) can be solved for $v_x^{n+1/2}(i,j,k)$ to yield the following:

$$v_x^{n+1/2}(i,j,k) = v_x^{n-1/2}(i,j,k) - \frac{2}{\rho c}\frac{c\Delta t}{\delta}\left[p^n(i+1,j,k) - p^n(i,j,k)\right]. \tag{14}$$

Note that the pressure node $p^n(i+1,j,k)$ will be zero since it is beyond the boundary and that the only difference between Eqs. (14) and (10) is a factor of two. This type of modification of the coefficient in some of the velocity update equations constitutes the heart of the conformal methods presented here.

In general, the true pressure-release boundary will not be perpendicular to one of the coordinates nor will it exactly bisect a cell. Thus, two obvious question arises: (1) When should a cell be split and (2) are there other cell-splitting schemes, besides a simple halving of the mass, that will improve the fidelity of a solution? The investigation of these questions has led us to develop three cell-splitting techniques as described in the following subsections. Results obtained using two of these methods are presented in the Sec. 2.

1.1 Quantized Split-Cell Method

In the quantized split-cell method, the mass of a cell associated with a velocity node is either that of the complete cell or that of half the cell. For example, a v_x node would either be updated using the usual update equation of (10) or the split-cell update equation of (14). The split-cell update equation is used if the pressure-release boundary intersects the midline of the velocity cell (i.e., the imaginary line on which the velocity node is located that passes through the middle of the cell) anywhere in the range $1/4 < L < 3/4$. The curvature and tilt of the boundary within a single cell are ignored. Thus, the decision to split the cell is based solely on where the boundary intersects the midline.

Pressure nodes are set to zero if any part of the pressure-release region is closer than one quarter of a spatial step size in any direction. Thus, for the pressure node $p(i, j, k)$ the neighboring points $([i \pm 1/4]\Delta x, j\Delta y, k\Delta z)$, $(i\Delta x, [j \pm 1/4]\Delta y, k\Delta z)$, and $(i\Delta x, j\Delta y, [k \pm 1/4]\Delta z)$ are checked and if any lie within the pressure-release region the pressure node is set to zero. When a velocity node is split, one of the neighboring pressure nodes will always be set to zero.

The quantized split-cell method can be run at the Courant limit (despite the fact it appears to contain cells at the boundaries that are half the size of the cells away from the boundary). This statement is based on numerous simulations employing a wide range resonant structures. Resonators are typically quick to show any numeric instabilities in an algorithm since, unlike in scattering problems, there is no means by which energy can exit the grid.

Depending on the implementation, there can be almost no computational cost associated with the quantized split-cell method. Some geometry preprocessing is required to determine which nodes are split and which are not. Because of the simple criterion used to determine which cells are split, the preprocessing is an insignificant portion of the total computation time (especially for three-dimensional objects). If the FDTD implementation uses an array to store the coefficients for each velocity node update equation, some coefficients will be that of the full cell and some will be that of the split cell. At this point time-stepping proceeds as usual and there is no additional cost incurred by the cell-splitting scheme. Alternatively, if coefficients arrays are not used (and they typically would not be if the bulk of the computational domain consists of a homogeneous material), the geometry preprocessing is used to construct a list of cells that are split. At each time step all nodes are updated in the usual way and then an additional subroutine (or function) is used to correct the fields at the boundary in accordance with the cell splitting. Since the list of boundary cells is associated with an object that is one dimension smaller than the total computational domain, the computational effort associated with this correction is small compared to the effort needed to advance the fields throughout the entire computational domain.

The quantized split-cell method is similar in many respects to the cell-splitting method that has been used to model perfect electrically conducting (PEC) boundaries [16] (see also

[20] and the citations therein). However, the electromagnetic cell-splitting method requires a global distortion of the grid to ensure that the PEC boundary is either aligned with the grid or passes through the diagonal of a cell. This distortion is necessitated by the fact that two vector fields are being advanced. In the acoustic FDTD method, where one scalar and one vector field are used, no grid distortions are required.

1.2 Hybrid Split-Cell Method

The hybrid split-cell method is similar to the quantized split-cell method. The criterion used for setting pressure nodes to zero is unchanged and the simulation can still be run at the Courant limit. As before, a velocity node is split if the pressure-release boundary intersects the cell's midline anywhere in the range $1/4 < L < 3/4$. However, when a cell is split, the mass is not simply cut in half. Instead, the split incorporates more geometry information by reducing the mass in accordance with the location of the intersection. The coefficient for a velocity cell that has been split is modified by a factor of $1/L$ rather than simply a factor of two. We identify this scheme as "hybrid" since it incorporates continuous geometry over a given range of intersections, but clips the grid to the staircase representation for intersections that fall outside the acceptable range.

1.3 Continuous Split-Cell Method

The continuous split-cell method is similar in many respects to the hybrid split-cell method. However, in this method the mass associated with a velocity cell is always dictated by the intersection of the pressure-release boundary with the midline of the cell. Thus the coefficients for split velocity cells are modified by $1/L$ regardless of L. Naturally, small values of L lead to large coefficients and numerical problems. To alleviate these problems, a threshold can be defined that L must exceed (or which $1 - L$ must not exceed) to result in splitting. When the threshold is one quarter of a cell, the hybrid method is obtained. (Pressure nodes can either be set to zero if they lie in the pressure-release region or be set to zero if any portion of some area around the node, as dictated by the threshold, falls within the pressure-release region.) For thresholds less than a quarter of cell, instabilities may occur. However, a stable solution can be obtained by reducing the Courant number.

The continuous method is similar to the conformal FDTD method recently presented by Dey and Mittra [17] to solve for electromagnetic scattering from PEC surfaces. An analytic (as opposed to numeric) analysis of the accuracy of this method can be found in [21]. Despite the success reported by Dey and Mittra for the electromagnetic case, it was found that there is no compelling reason to use such a method for the acoustic case. Despite the fact that the continuous method seems to include more geometry information when the thresholds are small, the results were typically no better than those obtained using the methods described above and, because of the reduced Courant number needed to obtain a stable solution, it takes longer to compute. Therefore, the results obtained using the continuous scheme will not be presented in the next section. (We have found no situation where the continuous scheme provides results consistently superior to the other cell-splitting methods.)

2 Results

Before presenting the results for the scatterer and resonator, we note that the following tests were designed to establish the accuracy of the conformal methods. To this end, the tests employ simple canonical geometries for which analytic solutions are available. Thus, we do not consider more "interesting" (and physically realistic) problems such as the scattering from the ocean surface. Furthermore, the code used to solve these problems was not optimized. For example, the solutions to these problems could easily be written to exploit the inherent symmetry, but this was not done. Hence we do not discuss the computer resources needed to solve these problems (we merely note that the simulations required only modest resources and were run on typically configured personal computers running the Linux operating system).

2.1 Scattering from a Sphere

Consider a spherical pressure-release region of radius a that is ensonified by a plane wave propagating in the z direction. A simple staircase realization of a sphere is shown in Fig. 3. The radius of the sphere is eight cells and this radius is used in all the subsequent calculations. A wire-frame box is used to indicate the size of the computational domain which was 39 by 39 by 39 cells and a heavy solid line shows the line over which scattered fields are recorded. The dashed lines indicate the plane over which the cross-sectional view of Fig. 4 is taken (both figures are drawn to scale). Ensonification is a Ricker wavelet with its peak spectral content corresponding to 20 cells per wavelength. The incident wave is introduced using a total-field/scattered-field formulation [22, 23, 24]. The simulation is run for 512 time steps and spectral information is obtained by taking the discrete Fourier transforms at the desired frequencies [25] and over the evaluation points. In this way broadband information is obtained from a single simulation. The grid is terminated using a fourth-order complementary operators method (COM) absorbing boundary condition [26, 27].

The exact solution for the pressure scattered by a sphere of radius a having a Dirichlet boundary is given by [28]

$$p_s(r,\theta,\phi) = -\sum_m (2m+1) i^{m+1} e^{iT_m(ka)} \sin(T(ka)) P_m(\cos(\theta)) h_m(kr) \tag{15}$$

where k is the wavenumber, $T_m(z) = \tan^{-1}(j_m(z)/n_m(z))$, P_m is the Legendre polynomial, and $e^{-i\omega t}$ dependence is understood. The spherical Bessel functions are given by $j_m(z) = \sqrt{\pi/2z} J_{m+1/2}(z)$, $n_m(z) = \sqrt{\pi/2z} N_{m+1/2}(z)$, and $h_m(z) = j_m(z) + i n_m(z)$.

All simulations were done using unitless quantities such as grid coordinates, cells per wavelength, and time steps. Thus the results can be scaled to any size and material (the only constraint is the size of the scatterer relative to a wavelength). However, to make the results more tangible, it is helpful to assign convenient physical dimensions to the results. Thus, we assume the sphere has a radius a of 1.0 m and the background material is water with a sound speed of 1500 m/s and a density of 1000 kg/m^3. Since the radius of the sphere spans eight cells, the corresponding spatial step size is 1/8 m and the length of one side of the computational domain is 39/8 m. The simulations are run at 98 percent of the Courant limit which corresponds to a temporal step size of 47.3 μs. Thus, the total duration of the simulations is 24.2 ms (512 time steps). (The simulations were not run exactly at the Courant limit in order to ensure the stability of the absorbing boundary condition. As mentioned previously, the cell-splitting method can be run at the limit.)

Figure 5 shows the scattered pressure measured over the evaluation line when the sphere is realized using a staircase approximation. Results are shown for a nominal discretizations of 10, 20, and 40 points per wavelength (PPW) corresponding to wavelengths of 1.25, 2.50, and 5.0 m. However, because of the duration of the simulation, the temporal step size, and the inherent resolution of the discrete Fourier transform, the discretizations that could be recorded directly from the FDTD simulation and that were closest to the nominal values had PPW's of 10.027, 19.386, and 41.542. Rather than interpolating the FDTD results to the nominal values, the true values were used in the exact solution. Nevertheless, the nominal values will be used to identify the results. Figure 5 shows the exact solution as well the solutions obtained using either the Stair-v or Stair-p realizations of the pressure-release surface. Clearly the Stair-p results are inferior to the Stair-v results.

Figure 6 shows the scattered pressure measured over the evaluation line when the pressure-release sphere is realized using either the quantized or the hybrid techniques. At discretization of 20 and 40 PPW, the split-cell results are difficult to distinguish from the exact solution. We observe that the hybrid method typically performs slightly better (but only slightly better) than the simpler quantized method.

Figure 7 shows the relative error, given by |FDTD − exact|/exact, for discretizations of 10 and 40 PPW for the Stair-v and hybrid methods. At 10 PPW, there are points where the Stair-v method has lower error than the hybrid method. However, the integrated error is obviously much lower for the hybrid method than for the staircased one. At 40 PPW, the hybrid results are superior to the Stair-v results for all points along the evaluation line.

2.2 Resonances of a Sphere

The resonant modes of a sphere of radius a that satisfies a Dirichlet (pressure-release) boundary condition are determined using the FDTD method. The realization of the sphere is similar to that shown in Figs. 3 and 4 except that the region of interest is interior to the sphere. The radius of the sphere is eight cells and will again, for the sake of concreteness, be assigned a physical dimension of 1.0 m.

The resonant modes are determined by introducing energy into the sphere, running the FDTD simulation while recording the temporal signal at an observation point, and then taking the (fast) Fourier transform (FFT) of the temporal signal and noting the resonances (the way in which energy is introduced "transparently" into the sphere can be found in [29]). The subsequent results were obtained by an automated process, i.e., a computer script was used to coordinate the tasks handled by separate programs and the resonances were found by the computer using a simple set of search rules. Simulations were run for 32,768 time steps at 98 percent of the Courant limit. The spectral resolution for this number of time steps, i.e., the difference between neighboring frequencies "bins" after taking the FFT, was 0.67 Hz. For consistency, the results presented here use the same Courant number as was used for the scattering problem. However, since the resonator does not require absorbing boundary conditions, simulations were also performed at the Courant limit and no unstable behavior was observed.

For a sphere of radius a the frequencies of the resonant modes are given by

$$f = \frac{c}{2\pi a}\zeta_{mp} \tag{16}$$

where ζ_{mp} represents the p^{th} zero of the spherical Bessel function of order m. The modes with the four lowest resonant frequencies correspond to values of m and p of (0,1), (1,1), (1,2), and (0,2). Table 1 lists these frequencies, in Hertz, for a sphere with a radius of 1.0 m. Also shown in this table are the resonant frequencies as determined from the FDTD simulations. The percent error, as given by $100 \times |\text{FDTD} - \text{exact}|/\text{exact}$ is given in parentheses. As was the case for the exterior problem, the Stair-v realization of the pressure-release surface is clearly better than the Stair-p realization. Furthermore, the split-cell methods yield results that are superior to those obtained using either of the staircase schemes. The hybrid method performs better than the quantized method at the lower frequencies, but does not provide a significant improvement at the higher frequencies.

3 Conclusions

The simplest FDTD implementation of a continuously varying pressure-release surface uses a staircase approximation of the surface. If the staircase approximation is constructed based solely on the locations of the pressure nodes themselves, significant errors may be introduced. An improved staircase representation is obtained by setting to zero those pressure nodes that have any of their neighboring velocities nodes beyond the pressure release boundary. This improvement yields results that are substantially better than those obtained from the method which is based solely on pressure-node locations. This improvement is realized despite the fact that this velocity-based method still ultimately produces a staircase representation of the surface and despite there being essentially no additional computational cost.

Further improvements can be realized by splitting velocity nodes adjacent to the pressure-release boundary. The splitting can be either a simple bisection (quantized split) or one that incorporates more information about where the boundary passes through the cell (hybrid split). The computational cost of these improvements are insignificant compared to the overall computational cost of the simulation. The results obtained using either of these splitting techniques are, in general, superior to those of either of the staircase methods. These methods permit the accurate solution of problems using a coarser grid than would otherwise be possible. For example, we envision that the split-cell methods will make possible the accurate determination in the time domain of scattering from large randomly rough surfaces such as the air-sea interface.

Acknowledgments

This work was supported by the Office of Naval Research, Code 321OA. We thank Todd Hefner of the Washington State University Physics Department for thought provoking discussions.

References

[1] K. S. Yee. Numerical solution of initial boundary value problems involving Maxwell's equations in isotropic media. *IEEE Trans. Antennas Propagat.*, 14(3):302–307, March 1966.

[2] R. Madariaga. Dynamics of an expanding circular fault. *Bull. Seism. Soc. Am.*, 66(3):639–666, June 1976.

[3] J. Virieux. *SH*-wave propagation in heterogeneous media: Velocity-stress finite difference method. *Geophysics*, 49(11):1933–1942, November 1984.

[4] J. Virieux. *P-SV* wave propagation in heterogeneous media: Velocity-stress finite difference method. *Geophysics*, 51(4):889–901, April 1986.

[5] C. J. Randall. Absorbing boundary condition for the elastic wave equation: Velocity-stress formulation. *Geophysics*, 54(9):1141–1152, September 1989.

[6] D. Botteldooren. Finite-difference time-domain simulation of low-frequency room acoustic problems. *J. Acoust. Soc. Am.*, 98(6):3302–3308, December 1995.

[7] J. LoVetri, D. Mardare, and G. Soulodre. Modeling of the seat dip effect using the finite-difference time-domain method. *J. Acoust. Soc. Am.*, 100(4):2204–2212, October 1996.

[8] R. A. Stephen. Modeling sea surface scattering by the time-domain finite-difference method. *J. Acoust. Soc. Am.*, 100(4):2070–2078, 1996.

[9] F. D. Hastings, J. B. Schneider, and S. L. Broschat. A finite-difference time-domain solution to scattering from a rough pressure-release surface. *J. Acoust. Soc. Am.*, 102(6):3394–3400, December 1997.

[10] J. G. Maloney and K. E. Cummings. Adaptation of FDTD techniques to acoustic modeling. In *11th Annual Review of Progress in Applied Computational Electromagnetics*, volume 2, pages 724–731, Monterey, CA, March 1995.

[11] X. Yuan, D. Borup, J. W. Wiskin, M. Berggren, R. Eidens, and S. A. Johnson. Formulation and validation of Berenger's PML absorbing boundary for the FDTD simulation of acoustic scattering. *IEEE Trans. Ultrasonics, Ferroelectrics, and Frequency Control*, 44(4):816–822, July 1997.

[12] Q. H. Liu and J. Tao. The perfectly matched layer for acoustic waves in absorptive media. *J. Acoust. Soc. Am.*, 102(4):2072–2082, October 1997.

[13] J. G. Maloney, K. L. Shlager, and G. S. Smith. A simple FDTD model for transient excitation of antennas by transmission lines. *IEEE Trans. Antennas Propagat.*, 42(2):289–292, February 1994.

[14] D. Botteldooren. Acoustical finite-difference time-domain simulation in a quasi-Cartesian grid. *J. Acoust. Soc. Am.*, 95(5):2313–2319, May 1994.

[15] T. G. Jurgens, A. Taflove, K. Umashankar, and T. G. Moore. Finite-difference time-domain modeling of curved surfaces. *IEEE Trans. Antennas Propagat.*, 40(4):357–366, April 1992.

[16] P. Mezzanotte, L. Roselli, and R. Sorrentino. A simple way to model curved metal boundaries in FDTD algorithm avoiding staircase approximation. *IEEE Microwave Guided Wave Lett.*, 5(8):267–269, August 1995.

[17] S. Dey and R. Mittra. A locally conformal finite-difference time-domain (FDTD) algorithm for modeling three-dimensional perfectly conducting objects. *IEEE Microwave Guided Wave Lett.*, 7(9):273–275, September 1997.

[18] F. D. Hastings, J. B. Schneider, S. L. Broschat, and E. I. Thorsos. Scattering from rough fluid-fluid interfaces using the finite-difference time-domain method. In *J. Acoust. Soc. Am.*, volume 101(5), State College, PA, June 1997.

[19] W. C. Chew. Electromagnetic theory on a lattice. *J. Appl. Phys.*, 75(10):4843–4850, May 1994.

[20] T. Weiland. Comment on "A simple way to model curved metal boundaries in FDTD algorithm avoiding staircase approximation". *IEEE Microwave Guided Wave Lett.*, 6(4):183–184, April 1996.

[21] J. B. Schneider and C. L. Wagner. Analytic analysis of the CP-FDTD and C-FDTD methods for offset planar boundaries. In *IEEE Antennas and Propagat. Soc. Int. Symposium*, volume 4, pages 1816–1819, Atlanta, GA, June 1998.

[22] D. E. Merewether, R. Fisher, and F. W. Smith. On implementing a numeric Huygen's source scheme in a finite difference program to illuminate scattering bodies. *IEEE Trans. Nucl. Sci.*, 27(6):1829–1833, December 1980.

[23] A. Taflove and K. Umashankar. Radar cross section of general three-dimensional scatterers. *IEEE Trans. Electromagn. Compat.*, EMC-25(4):433–440, November 1983.

[24] J. Fang. *Time Domain Finite Difference Computation for Maxwell's Equations.* PhD thesis, University of California at Berkeley, Berkeley, CA, 1989.

[25] C. M. Furse, S. P. Mathur, and O. P. Gandhi. Improvements to the finite-difference time-domain method for calculating the radar cross section of a perfectly conducting target. *IEEE Trans. Microwave Theory Tech.*, 38(7):919–927, July 1990.

[26] O. M. Ramahi. Complementary boundary operators for wave propagation problems. *J. Comput. Phys.*, 133:113–128, 1997.

[27] J. B. Schneider and O. M. Ramahi. The complementary operators method applied to acoustic finite-difference time-domain simulations. *J. Acoust. Soc. Am.*, 104(2):686–693, August 1998.

[28] P. M. Morse and H. Feshbach. *Methods of Theoretical Physics.* McGraw-Hill, New York, NY, 1953.

[29] J. B. Schneider, C. L. Wagner, and S. L. Broschat. Implementation of transparent sources embedded in acoustic finite-difference time-domain grids. *J. Acous. Soc. Am.*, 103(1):136–142, January 1998.

	Resonant Mode			
	1	2	3	4
Exact	750.0	1072	1375	1500
Stair-p	724.0 (3.45)	1033 (3.65)	1318 (4.17)	1439 (4.01)
Stair-v	760.2 (1.36)	1085 (1.16)	1384 (0.61)	1510 (0.71)
Quantized	755.6 (0.76)	1078 (0.49)	1374 (0.09)	1502 (0.16)
Hybrid	752.4 (0.33)	1074 (0.14)	1367 (0.61)	1496 (0.23)

Table 1: The four lowest resonant frequencies, in Hertz, of a sphere with a radius of 1.0 m bounded by a pressure-release region. The exact frequencies are given in the top row. The data in the other rows was obtained using the FDTD method and the percent error is given in parentheses.

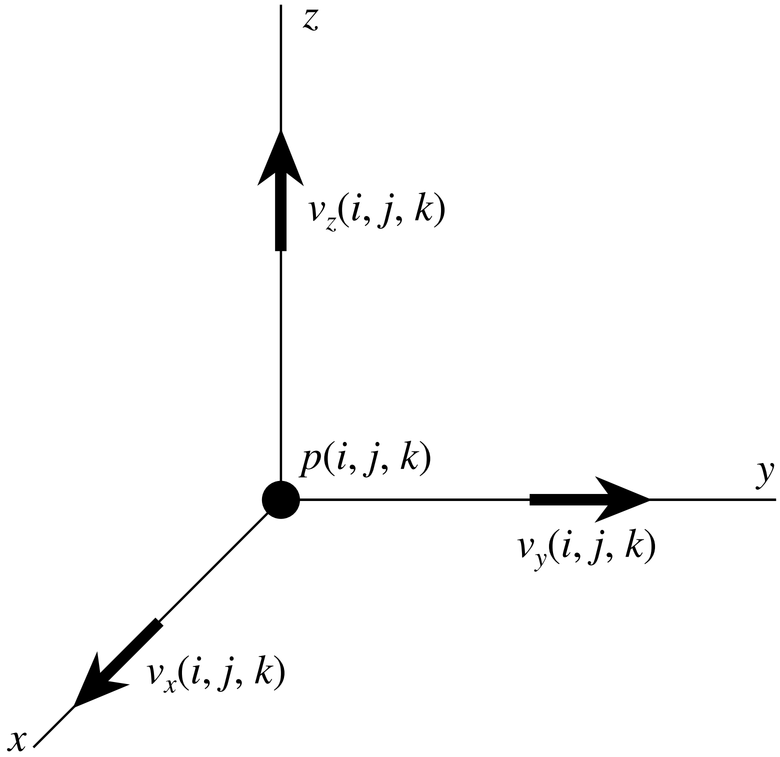

Figure 1: The location of the velocity nodes relative to a pressure node for a given set of spatial indices. The velocity and pressure nodes are also offset in time by half of the temporal step size.

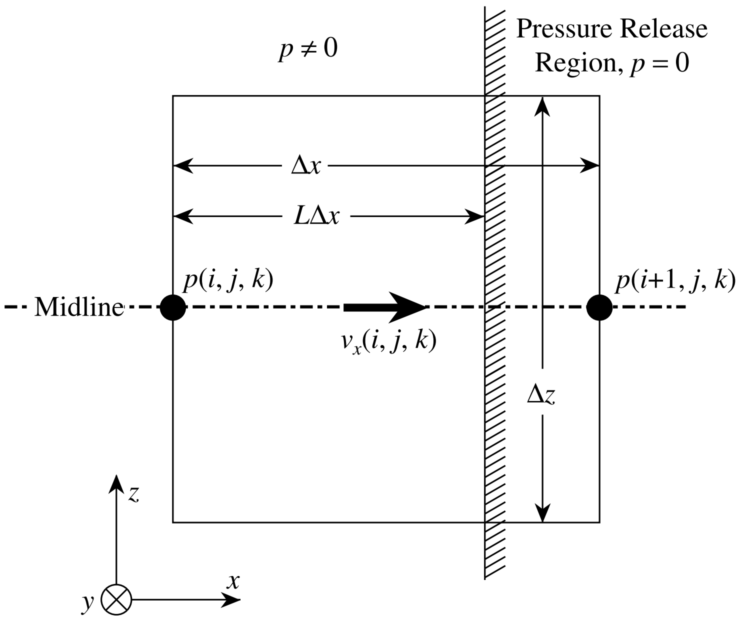

Figure 2: Cross section of the cell associated with a v_x node.

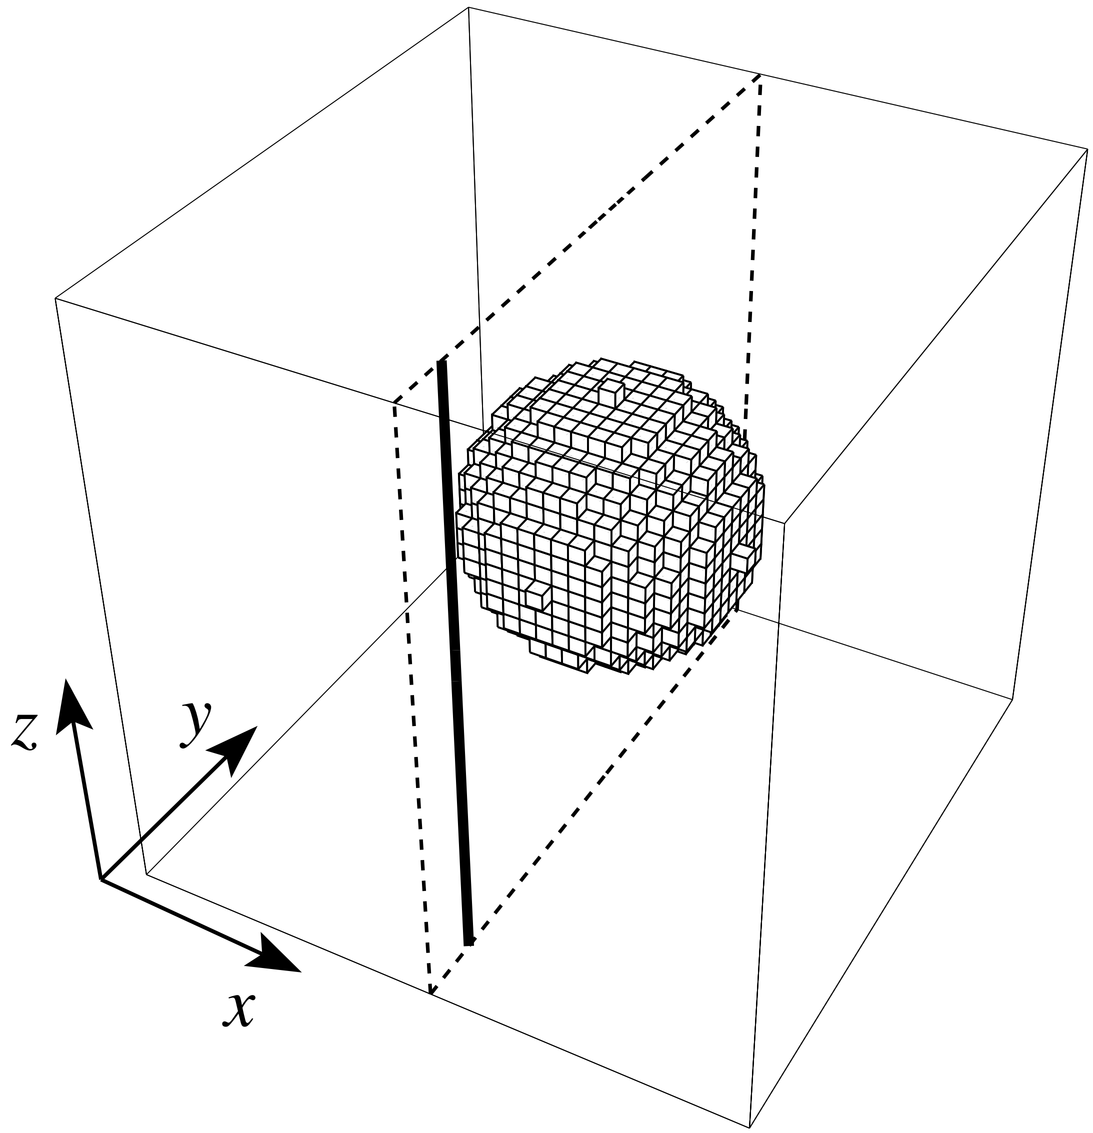

Figure 3: Depiction of the staircase FDTD realization of a spherical scatterer. The sphere has a radius of eight cells and the entire computational domain is 39 cells to a side. The heavy vertical line indicates the line over which the fields are recorded. The dashed lines indicate the plane over which the following figure is based.

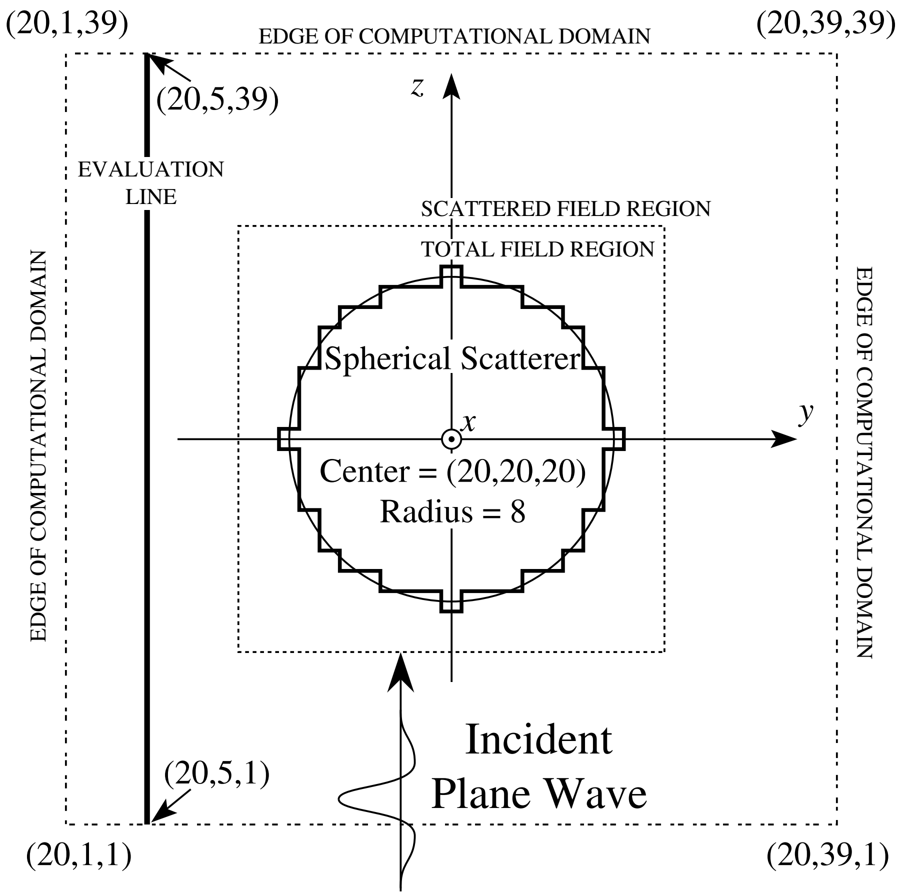

Figure 4: Cross-sectional view of the computational domain. The drawing is to scale. The coordinates shown are grid locations (spatial indices). Subsequent results assume the origin is at the center of the grid and use physical dimensions rather than grid coordinates.

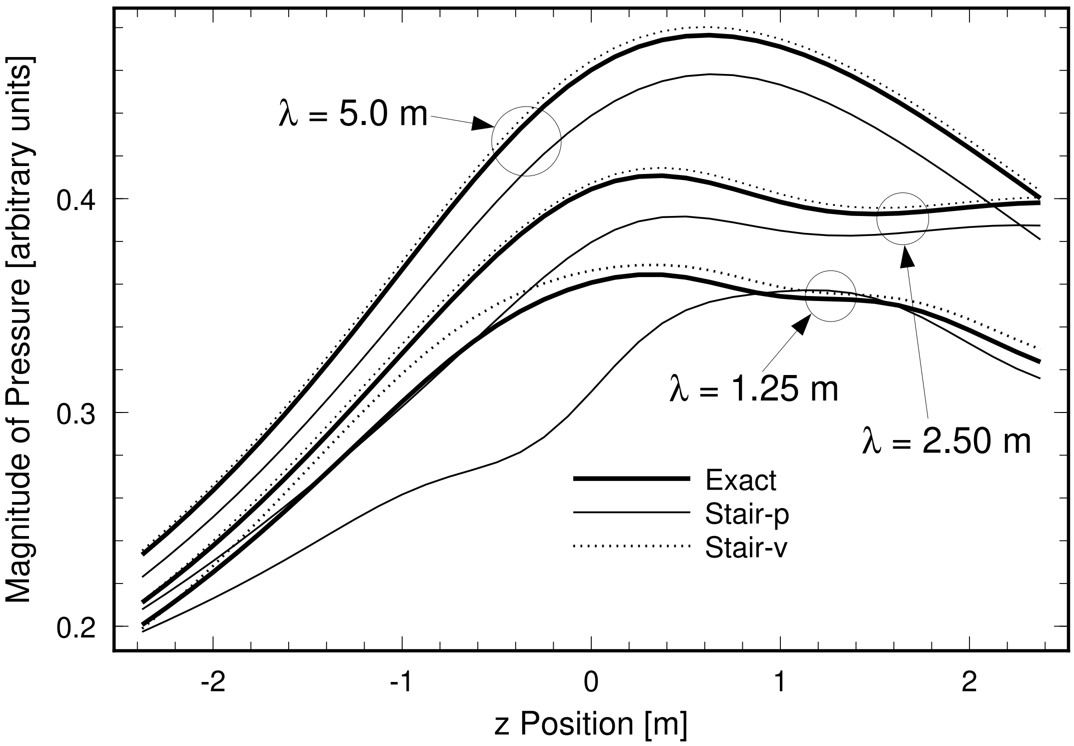

Figure 5: Scattered pressure at evaluation line obtained using a staircase realization of the sphere.

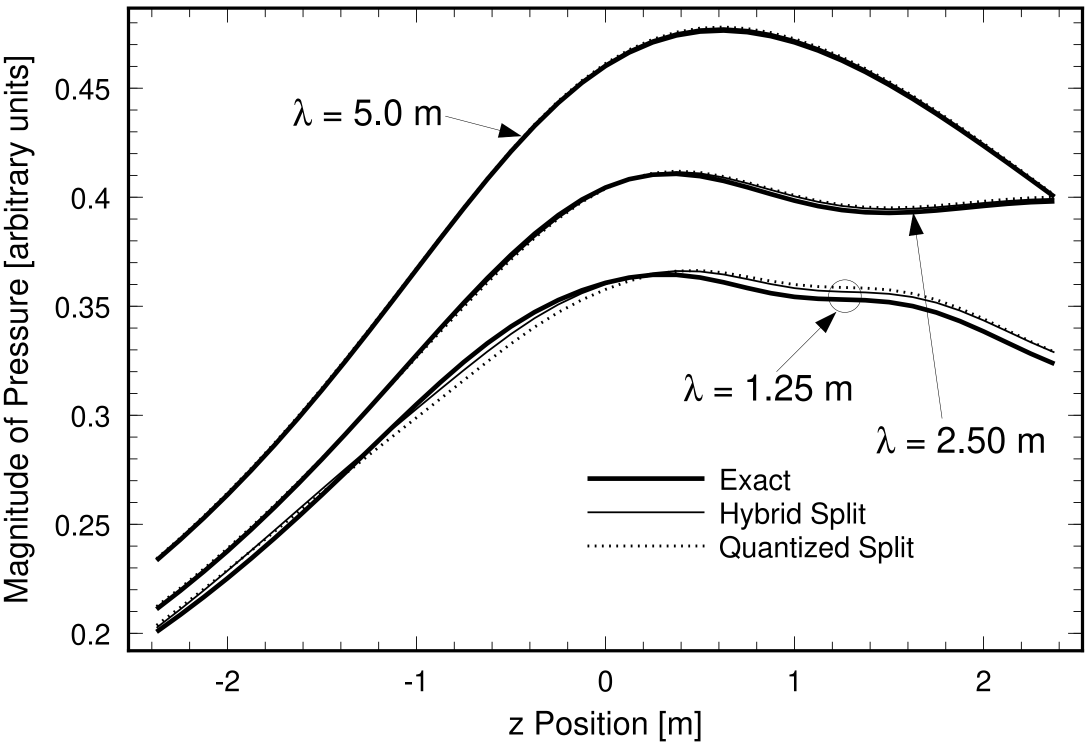

Figure 6: Scattered pressure at evaluation line obtained using a split-cell realization of the sphere.

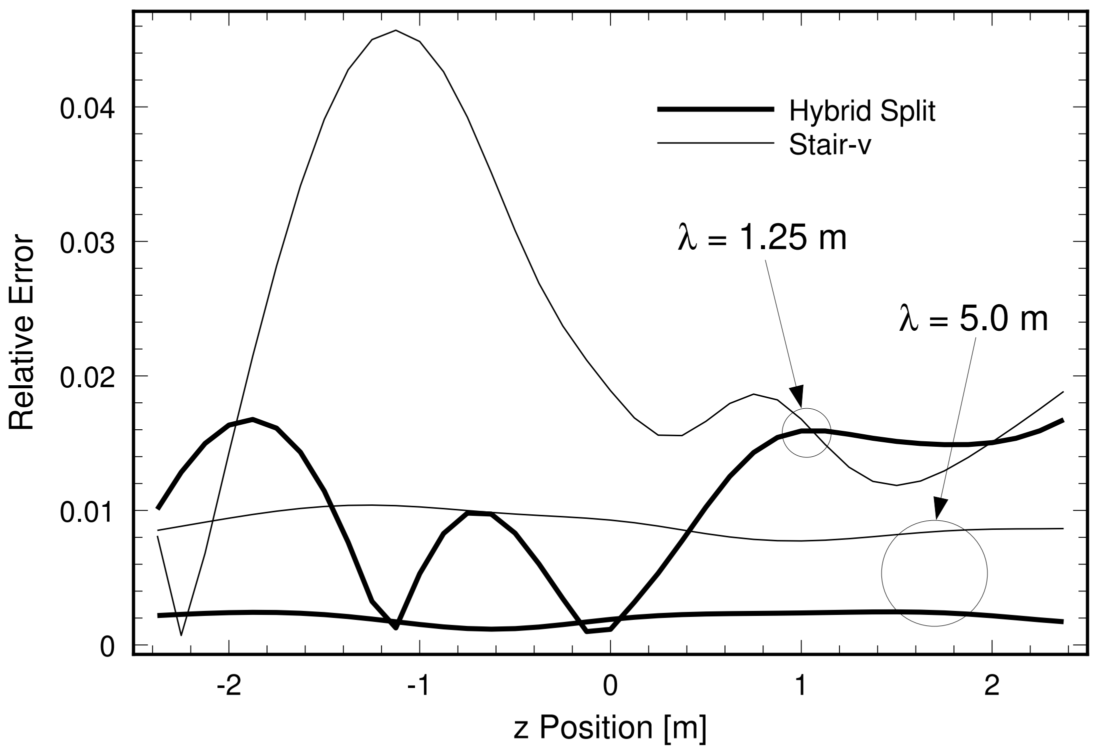

Figure 7: Relative error in the scattered field obtained using the Stair-v or hybrid split-cell method.