A k -space method for acoustic propagation using coupled first-order equations in three dimensions

Jason C. Tillett

Department of Electrical and Computer Engineering, University of Rochester, Rochester, New York 14627

Mohammad I. Daoud

Department of Electrical and Computer Engineering and Robarts Research Institute, University of Western Ontario, London, Ontario N6A 5B9, Canada

James C. Lacefield

Department of Electrical and Computer Engineering, Department of Medical Biophysics, and Robarts Research Institute, University of Western Ontario, London, Ontario N6A 5B9, Canada

Robert C. Waag

Department of Electrical and Computer Engineering and Department of Imaging Sciences, University of Rochester, Rochester, New York 14627

(Received 31 December 2008; revised 26 May 2009; accepted 28 May 2009)

A previously described two-dimensional k -space method for large-scale calculation of acoustic wave propagation in tissues is extended to three dimensions. The three-dimensional method contains all of the two-dimensional method features that allow accurate and stable calculation of propagation. These features are spectral calculation of spatial derivatives, temporal correction that produces exact propagation in a homogeneous medium, staggered spatial and temporal grids, and a perfectly matched boundary layer. Spectral evaluation of spatial derivatives is accomplished using a fast Fourier transform in three dimensions. This computational bottleneck requires all-to-all communication; execution time in a parallel implementation is therefore sensitive to node interconnect latency and bandwidth. Accuracy of the three-dimensional method is evaluated through comparisons with exact solutions for media having spherical inhomogeneities. Large-scale calculations in three dimensions were performed by distributing the nearly 50 variables per voxel that are used to implement the method over a cluster of computers. Two computer clusters used to evaluate method accuracy are compared. Comparisons of k -space calculations with exact methods including absorption highlight the need to model accurately the medium dispersion relationships, especially in large-scale media. Accurately modeled media allow the k -space method to calculate acoustic propagation in tissues over hundreds of wavelengths.

© 2009 Acoustical Society of America. [DOI: 10.1121/1.3158857]

PACS number(s): 43.58.Ta, 43.20.Fn, 43.80.Qf, 43.35.Fj [TDM]

Pages: 1231–1244

I. INTRODUCTION

The calculation of acoustic propagation through tissue over a large number of wavelengths is a challenging problem because the resources required for computations over such distances can exceed commonly available desktop facilities when the computations are performed in three dimensions. The solution of this problem is an integral part of inversion algorithms and a valuable tool for studying aberration estimation and correction techniques in ultrasound imaging. Furthermore, large-scale calculations can model research and clinical experimental configurations at full scale. The problem, sometimes called the forward problem, is solved in this paper by extending a computationally efficient solution to three dimensions from two dimensions and by using a cluster of computers with a message-passing interface (MPI). The efficient solution is an extension of the so-called k -space method described in the work of Tabei *et al.*¹ In that work, which extended second-order formulations^{2–5} of this method, coupled first-order equations were used, a perfectly matched layer (PML) (Ref. 6) was incorporated on the boundary of

the computational domain, and loss was described by relaxation absorption,⁷ which is the dominant cause of loss in tissue. The k -space method was originally described by Bojarski.^{8,9}

Although k -space code that calculates propagation in three dimensions can be executed on a typical desktop computer, the computational scale is limited by the physical memory that is available because the number of cells or voxels must be less than 9×10^7 for a desktop computer with 16 Gbytes of memory. This translates for the current method into a $416 \times 416 \times 416$ grid for a representative 45 single-precision floating-point variables per grid cell. The large number of variables is comprised of, but not limited to, a decomposed pressure, velocity field components, medium parameters (one of which is required on four spatial grids), relaxation times and generalized compressibilities (one for each absorption process), and parameters of the PML. Using a minimum of four points-per-wavelength at the highest frequency contained in a typical pulse propagating through the medium, desktop computing resources fall short of applica-

FIG. 1. Time required to execute a three-dimensional FFT on a single computer compared to the time to execute a three-dimensional FFT on a cluster of 32 computers for a grid size of $N=64, 128, 256$, or 512 . Each point is the average of ten trials. The single-computer architecture is a 64-bit AMD Opteron having 16 Gbytes of random access memory (RAM) running Windows XP x64 and a 64-bit version of MATLAB. The cluster of computers consists of Pentium-III 1.4 GHz processors each having 0.5 Gbyte of RAM, running Linux, and connected via a 1-Gbit/s Ethernet switched network.

tions such as full-scale calculation of scattering in the breast, evaluation of aberration correction techniques, and simulation of b -scan imaging.¹⁰ The three-dimensional k -space code described here for large-scale calculations is designed to run on relatively low-cost Beowulf clusters¹¹ using MPI, where the amount of physical memory is the aggregate of all of the individual memories, and uses a slab-decomposition of the medium and associated variables, where each compute node handles a range of indices on one of the three-dimensions. Memory capacity is not the only limiting factor. Execution time as shown in Fig. 1 increases as a power of the size of the three-dimensional fast Fourier transform (FFT), the numerical method of the three-dimensional k -space code containing the bulk of the necessary arithmetic operations required to implement the algorithm. However, if the three-dimensional FFT is distributed over many nodes, as can be accomplished using the fastest-Fourier-transform-in-the-west (FFTW) software,¹² the exponent of the power-law scaling is effectively decreased.

Second-order formulations of the k -space method have been extended to three dimensions.^{2,13} These formulations can have absorbing boundary conditions and absorption. However, the PML incorporated in the first-order implementation can suppress reverberations more efficiently for thinner boundary layers. Additionally, relaxation absorption included in the first-order implementation allows for a frequency dependence like that observed in tissues. Other k -space methods for wave propagation in three dimensions are available^{14,15} but constrain the incident field and are not amenable to incorporation of relaxation absorption that is dominant in tissues.

Analysis of errors and performance is of particular interest in the three-dimensional k -space forward-problem solution because high-cost specialized computing platforms, such as clusters of computers, must be used for the large-scale

problems of interest. Specifically, any increase in the number of points-per-wavelength required to achieve a specified accuracy translates directly into larger grid sizes.

The k -space method is ideally suited for calculating propagation in media where fluctuations in compressibility and density are small as they are in soft tissue. In the calculation, the scattering medium is sampled but previous experience has shown that sudden changes in medium parameters introduce artifacts. Such artifacts have been addressed by using half-band filtering¹⁶ that has been shown to achieve a good balance between suppressing artifacts and maintaining high spatial frequencies. This balance results in scattered wave amplitudes that closely match results of exact calculations. Accurate Fourier decomposition of the medium produces better half-band filtering but typically requires oversampling the medium and, thus, increased k -space grid size. Since k -space grid sizes for large-scale computations are already large, oversampling can become a problem even for a large cluster. Therefore, attention is given to characterize the impact of varying amounts of oversampling.

The treatment in this paper begins with a mathematical formulation that explicitly extends the two-dimensional k -space method to three dimensions and compares the method to developments in related publications. The methods used to evaluate performance are described next. Results of numerical experiments are then presented. This is followed by a discussion of the results. Finally, conclusions are drawn from the reported studies.

II. MATHEMATICAL FORMULATION

A. Physical acoustic foundations

The development of equations for application of the k -space method to calculate acoustic wave propagation in liquids over a scale of many wavelengths in three dimensions is summarized here and related to results in other publications. Throughout the summary, the notation and conventions are consistent. The use of different notations and conventions from publication to publication is noted.

The summary begins with the equation for conservation of mass

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{v}) = 0, \quad (1)$$

the equation for momentum time rate of change

$$\rho \left(\frac{\partial}{\partial t} + \mathbf{v} \cdot \nabla \right) \mathbf{v} = -\nabla p, \quad (2)$$

and the equation for conservation of energy neglecting thermal conduction

$$\rho \frac{D}{Dt} \left(\frac{1}{2} v^2 + u \right) = -\nabla \cdot (p \mathbf{v}), \quad (3)$$

which are commonly developed in acoustic textbooks, e.g., Ref. 17 or Ref. 18. In these equations, ρ is medium density, $\mathbf{v}$ is velocity with $\|\mathbf{v}\|=v$, p is total pressure, u is internal energy per unit volume, and the operator D/Dt is the time rate of change evaluated in a coordinate system that is sta-

tionary with respect to a moving infinitesimal volume element and is equivalent to the operator in parentheses on the left side of Eq. (2). Relaxation absorption is included using the phenomenological theory of irreversible thermodynamics^{19,20} in which time rate of change in extensive parameters is related to deviations of the extensive parameters from equilibrium values and flows of energy and particles are driven by generalized forces (i.e., affinities) to obtain a Gibbs–Duhem relation²¹

$$T \frac{Ds}{Dt} = \frac{Du}{Dt} + p \frac{D(\rho^{-1})}{Dt} + \sum_{\nu} A_{\nu} \frac{Dn_{\nu}}{Dt}. \quad (4)$$

In this equation, T is temperature, s is entropy per unit volume, n_{ν} is particle number density, A_{ν} is particle affinity, and the subscript ν denotes the particle species.

Assuming the acoustic perturbations are small and neglecting effects of gravity, viscosity, and heat conduction lead Eqs. (1)–(4) as shown in Ref. 7 to a linear set of equations that are

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho_0 \mathbf{v}) = 0, \quad (5)$$

$$\rho_0 \frac{\partial \mathbf{v}}{\partial t} = -\nabla p, \quad (6)$$

$$\frac{1}{c^2} \frac{\partial p}{\partial t} - \left(\frac{\partial p}{\partial t} + \mathbf{v} \cdot \nabla \rho_0 \right) + \sum_{\nu} \rho_0 \kappa_{\nu} \frac{\partial (\Delta \xi_{\nu})}{\partial t} = 0, \quad (7)$$

and

$$\left(\frac{\partial}{\partial t} + \frac{1}{\tau_{\nu}} \right) (\Delta \xi_{\nu}) = -\frac{\partial p}{\partial t}. \quad (8)$$

In these equations, ρ_0 is the equilibrium density, c is acoustic wave speed, κ_{ν} and τ_{ν} are the isothermal compressibility and relaxation time constant, respectively, of species ν , and $\Delta \xi_{\nu}$ represents a pressure change produced by the deviation of n_{ν} from its equilibrium value $n_{\nu,0}$. This pressure change is given by

$$\Delta \xi_{\nu} = (n_{\nu} - n_{\nu,0}) / (\partial n_{\nu,0} / \partial p)_T. \quad (9)$$

Substitution of Eq. (9) in the above equations to eliminate $\Delta \xi_{\nu}$ yields equations for acoustic pressure and particle velocity that are the starting point in parallel developments of time-domain calculation algorithms found in Refs. 1 and 22. These equations, without terms incorporating a PML that will be included later, are

$$\rho_0(\mathbf{x}) \frac{\partial \mathbf{v}(\mathbf{x}, t)}{\partial t} = -\nabla p(\mathbf{x}, t) \quad (10)$$

and

$$\kappa(\mathbf{x}, t) \otimes \frac{\partial p(\mathbf{x}, t)}{\partial t} = -\nabla \cdot \mathbf{v}(\mathbf{x}, t). \quad (11)$$

Equations (10) and (11) are three-dimensional extensions of Eqs. (13)–(16) in Ref. 1 or, equivalently, Eqs. (7)–(10) in Ref. 22 that are written in each reference as components in two-dimensional space.

Expansion of Eqs. (10) and (11) in Cartesian coordinates yields

$$\begin{aligned} \frac{\partial}{\partial x} p(\mathbf{x}, t) &= -\rho_0(\mathbf{x}) \frac{\partial}{\partial t} v_x(\mathbf{x}, t), \\ \frac{\partial}{\partial y} p(\mathbf{x}, t) &= -\rho_0(\mathbf{x}) \frac{\partial}{\partial t} v_y(\mathbf{x}, t), \\ \frac{\partial}{\partial z} p(\mathbf{x}, t) &= -\rho_0(\mathbf{x}) \frac{\partial}{\partial t} v_z(\mathbf{x}, t), \end{aligned} \quad (12)$$

and

$$\begin{aligned} \kappa(\mathbf{x}, t) \otimes \frac{\partial p(\mathbf{x}, t)}{\partial t} + \frac{\partial}{\partial x} v_x(\mathbf{x}, t) + \frac{\partial}{\partial y} v_y(\mathbf{x}, t) + \frac{\partial}{\partial z} v_z(\mathbf{x}, t) \\ = 0. \end{aligned} \quad (13)$$

In these equations, $\otimes$ denotes temporal convolution and $\kappa(\mathbf{x}, t)$ is the generalized compressibility given by⁷

$$\kappa(\mathbf{x}, t) = \kappa_{\infty}(\mathbf{x}) \delta(t) + \sum_{\nu} \frac{\kappa_{\nu}(\mathbf{x})}{\tau_{\nu}(\mathbf{x})} e^{-t/\tau_{\nu}(\mathbf{x})} H(t), \quad (14)$$

where $H(t)$ is the Heaviside step function and

$$\kappa_{\infty}(\mathbf{x}) = \kappa_0(\mathbf{x}) - \sum_{\nu} \kappa_{\nu}(\mathbf{x}) \quad (15)$$

in which

$$\kappa_0(\mathbf{x}) = 1/[\rho(\mathbf{x})c^2(\mathbf{x})]. \quad (16)$$

B. Inclusion of a PML

A PML on the boundary of the computation is integrated into the development by assuming an $e^{j\omega t}$ time dependence, which is the same convention assumed in Ref. 23, and transforming Eqs. (12) and (13) into the temporal-frequency domain. The resulting equations are^a

$$\begin{aligned} \frac{\partial}{\partial x} p(\mathbf{x}, \omega) &= -j\omega \rho_0(\mathbf{x}) v_x(\mathbf{x}, \omega), \\ \frac{\partial}{\partial y} p(\mathbf{x}, \omega) &= -j\omega \rho_0(\mathbf{x}) v_y(\mathbf{x}, \omega), \\ \frac{\partial}{\partial z} p(\mathbf{x}, \omega) &= -j\omega \rho_0(\mathbf{x}) v_z(\mathbf{x}, \omega), \end{aligned} \quad (17)$$

and

$$\nabla \cdot \mathbf{v}(\mathbf{x}, \omega) = -j\omega \kappa(\mathbf{x}, \omega) p(\mathbf{x}, \omega). \quad (18)$$

Next, attenuation terms that are zero within the medium and nonzero in the PML are introduced. These terms are adjusted to attenuate waves in the boundary regions and eliminate reverberations of outgoing waves. An attenuation term $\alpha^*(\mathbf{x}, \omega)$ associated with generalized compressibility is assumed to have dispersion like that of the generalized compressibility. An attenuation term $\alpha(\mathbf{x})$ associated with the density is assumed to be dispersionless because the density is dispersionless. The asterisk notation is the same as the nota-

tion in Refs. 1 and 22 but is opposite the convention in Ref. 23. The modified equations are

$$\begin{aligned}\frac{\partial}{\partial x}p(\mathbf{x}, \omega) &= -j\omega\rho_0(\mathbf{x})v_x(\mathbf{x}, \omega) - \alpha(\mathbf{x})v_x(\mathbf{x}, \omega), \\ \frac{\partial}{\partial y}p(\mathbf{x}, \omega) &= -j\omega\rho_0(\mathbf{x})v_y(\mathbf{x}, \omega) - \alpha(\mathbf{x})v_y(\mathbf{x}, \omega), \\ \frac{\partial}{\partial z}p(\mathbf{x}, \omega) &= -j\omega\rho_0(\mathbf{x})v_z(\mathbf{x}, \omega) - \alpha(\mathbf{x})v_z(\mathbf{x}, \omega),\end{aligned}\quad (19)$$

and

$$\nabla \cdot \mathbf{v}(\mathbf{x}, \omega) = -j\omega\kappa(\mathbf{x}, \omega)p(\mathbf{x}, \omega) - \alpha^*(\mathbf{x}, \omega)p(\mathbf{x}, \omega). \quad (20)$$

Two steps are now taken. The first is to decompose pressure artificially into three Cartesian components by writing

$$p(\mathbf{x}, \omega) = p_x(\mathbf{x}, \omega) + p_y(\mathbf{x}, \omega) + p_z(\mathbf{x}, \omega). \quad (21)$$

This decomposition enables inclusion of the PML boundary conditions. The second is to assume the PML attenuation coefficients are anisotropic. Decomposing the pressure and separating Eq. (20) into three equations yield

$$\begin{aligned}\frac{\partial}{\partial x}v_x(\mathbf{x}, \omega) &= -j\omega\kappa(\mathbf{x}, \omega)p_x(\mathbf{x}, \omega) - \alpha_x^*(\mathbf{x}, \omega)p_x(\mathbf{x}, \omega), \\ \frac{\partial}{\partial y}v_y(\mathbf{x}, \omega) &= -j\omega\kappa(\mathbf{x}, \omega)p_y(\mathbf{x}, \omega) - \alpha_y^*(\mathbf{x}, \omega)p_y(\mathbf{x}, \omega), \\ \frac{\partial}{\partial z}v_z(\mathbf{x}, \omega) &= -j\omega\kappa(\mathbf{x}, \omega)p_z(\mathbf{x}, \omega) - \alpha_z^*(\mathbf{x}, \omega)p_z(\mathbf{x}, \omega).\end{aligned}\quad (22)$$

If α^* is anisotropic, then Eq. (22) becomes

$$\begin{aligned}\frac{\partial}{\partial x}v_x(\mathbf{x}, \omega) &= -j\omega\kappa(\mathbf{x}, \omega)p_x(\mathbf{x}, \omega) - \alpha_x^*(\mathbf{x}, \omega)p_x(\mathbf{x}, \omega), \\ \frac{\partial}{\partial y}v_y(\mathbf{x}, \omega) &= -j\omega\kappa(\mathbf{x}, \omega)p_y(\mathbf{x}, \omega) - \alpha_y^*(\mathbf{x}, \omega)p_y(\mathbf{x}, \omega), \\ \frac{\partial}{\partial z}v_z(\mathbf{x}, \omega) &= -j\omega\kappa(\mathbf{x}, \omega)p_z(\mathbf{x}, \omega) - \alpha_z^*(\mathbf{x}, \omega)p_z(\mathbf{x}, \omega).\end{aligned}\quad (23)$$

If α is anisotropic, then Eq. (19) becomes

$$\begin{aligned}\frac{\partial}{\partial x}p(\mathbf{x}, \omega) &= -j\omega\rho_0(\mathbf{x})v_x(\mathbf{x}, \omega) - \alpha_x(\mathbf{x})v_x(\mathbf{x}, \omega), \\ \frac{\partial}{\partial y}p(\mathbf{x}, \omega) &= -j\omega\rho_0(\mathbf{x})v_y(\mathbf{x}, \omega) - \alpha_y(\mathbf{x})v_y(\mathbf{x}, \omega), \\ \frac{\partial}{\partial z}p(\mathbf{x}, \omega) &= -j\omega\rho_0(\mathbf{x})v_z(\mathbf{x}, \omega) - \alpha_z(\mathbf{x})v_z(\mathbf{x}, \omega).\end{aligned}\quad (24)$$

Equations (23) and (24) can be further simplified using a relation between $\alpha^*(\mathbf{x}, \omega)$ and $\alpha(\mathbf{x})$ that can be derived by enforcing the constraint that normally incident plane waves at the boundaries of the medium are not reflected. This condition follows from assuming a plane-wave solution with the forms

$$\tilde{\mathbf{p}} = \tilde{\mathbf{p}}_0 e^{-j\mathbf{B} \cdot \mathbf{x}} \quad (25)$$

and

$$\tilde{\mathbf{v}} = \tilde{\mathbf{v}}_0 e^{-j\mathbf{B} \cdot \mathbf{x}}, \quad (26)$$

substituting Eq. (25) into Eq. (23) and Eq. (26) into Eq. (24), assuming κ and ρ are independent of $\mathbf{x}$, solving for the unknowns $\mathbf{B}$ and $\tilde{\mathbf{v}}$, and using the techniques detailed in Ref. 6. The condition is

$$\alpha_{(\cdot)}^*(\mathbf{x}, t) = \kappa(\mathbf{x}, t) \frac{\alpha_{(\cdot)}(\mathbf{x})}{\rho(\mathbf{x})}, \quad (27)$$

in which $(\cdot)$ denotes x , y , or z . When Eqs. (23) and (24) are Fourier transformed back to the time domain and reflections are eliminated using Eq. (27) the set of time-dependent first-order k -space equations that describe acoustic propagation in three dimensions in a medium that has relaxation absorption and a PML on the boundary becomes

$$\begin{aligned}\rho(\mathbf{x}) \left[\frac{\partial v_{(\cdot)}(\mathbf{x}, t)}{\partial t} + \alpha_{(\cdot)}(\mathbf{x})v_{(\cdot)}(\mathbf{x}, t) \right] \\ = - \frac{\partial [p_x(\mathbf{x}, t) + p_y(\mathbf{x}, t) + p_z(\mathbf{x}, t)]}{\partial(\cdot)}\end{aligned}\quad (28)$$

and

$$\kappa(\mathbf{x}, t) \otimes \left[\frac{\partial p_{(\cdot)}(\mathbf{x}, t)}{\partial t} + \alpha_{(\cdot)}(\mathbf{x})p_{(\cdot)}(\mathbf{x}, t) \right] = - \frac{\partial v_{(\cdot)}(\mathbf{x}, t)}{\partial(\cdot)}. \quad (29)$$

The above set of equations is an extension of Eqs. (13)–(16) of Ref. 1 from two dimensions to three dimensions. If κ and α^* are assumed to be independent of temporal frequency and Eqs. (23) and (24) are Fourier transformed back to the time domain, then, using

$$\alpha = \alpha^* \rho / \kappa, \quad (30)$$

a symmetric set of equations in three dimensions results. These equations are the three-dimensional analogs of the two-dimensional equations (11)–(14) in Ref. 23. If Eqs. (23) and (24) are Fourier transformed back to the time domain, Eqs. (7)–(10) in Ref. 22 result.

C. Solution for scattered pressure

Equations (28) and (29) describe the propagation of total pressure. Corresponding equations that describe the propagation of only the scattered pressure can be derived by substituting $p = p_s + p_i$ into Eqs. (28) and (29), noting the incident pressure satisfies the same set of equations in a homogeneous background medium, and eliminating the total pressure in the resulting set of equations. To present results with fewer subscripts in what follows, the subscript i is used to denote incident quantities and the absence of a subscript denotes scattered quantities. A compact set of equations, which is the solution for scattered pressure in continuous time and continuous space, is written by defining a state variable $S_p(\mathbf{x}, t)$ for each relaxation process. The state variables are filtered versions of the pressure field and are given by

$$S_\nu(\mathbf{x}, t) \equiv p(\mathbf{x}, t) \otimes \frac{\kappa_\nu(\mathbf{x})}{\tau_\nu(\mathbf{x})} e^{-t/\tau_\nu(\mathbf{x})} H(t). \quad (31)$$

These state variables combined with identities²² for convolutions including time derivatives allow the equations for scattered pressure to be written as

$$\begin{aligned} \frac{\partial v_{(\cdot)}(\mathbf{x}, t)}{\partial t} + \alpha_{(\cdot)}(\mathbf{x}) v_{(\cdot)}(\mathbf{x}, t) \\ = - \frac{1}{\rho(\mathbf{x})} \frac{\partial p(\mathbf{x}, t)}{\partial (\cdot)} - \left[\frac{1}{\rho(\mathbf{x})} + \frac{1}{\rho_0} \right] \frac{\partial p_i(\mathbf{x}, t)}{\partial (\cdot)} \end{aligned} \quad (32)$$

and

$$\begin{aligned} [\kappa_\infty(\mathbf{x})] \left[\frac{\partial p_{(\cdot)}(\mathbf{x}, t)}{\partial t} + \alpha_{(\cdot)}(\mathbf{x}) p_{(\cdot)}(\mathbf{x}, t) \right] + \sum_\nu \frac{\kappa_\nu(\mathbf{x})}{\tau_\nu(\mathbf{x})} p_{(\cdot)}(\mathbf{x}, t) \\ + \sum_\nu \left[-\frac{1}{\tau_\nu(\mathbf{x})} + \alpha_{(\cdot)}(\mathbf{x}) \right] S_\nu^{(\cdot)}(\mathbf{x}, t) \kappa_\nu(\mathbf{x}) \\ = - \frac{\partial v_{(\cdot)}(\mathbf{x}, t)}{\partial (\cdot)} - [\kappa_\infty(\mathbf{x}) - \kappa_0] \left[\frac{\partial p_{(\cdot),i}(\mathbf{x}, t)}{\partial t} \right. \\ \left. + \alpha_{(\cdot)}(\mathbf{x}) p_{(\cdot),i}(\mathbf{x}, t) \right] - \sum_\nu \frac{\kappa_\nu(\mathbf{x})}{\tau_\nu(\mathbf{x})} p_{(\cdot),i}(\mathbf{x}, t) \\ - \sum_\nu \left[-\frac{1}{\tau_\nu(\mathbf{x})} + \alpha_{(\cdot)}(\mathbf{x}) \right] S_\nu^{(\cdot),i}(\mathbf{x}, t) \kappa_\nu(\mathbf{x}). \end{aligned} \quad (33)$$

To facilitate a compact but more detailed presentation of these equations on a discrete spatial grid and at discrete temporal points, the following quantities are defined:

$$\mu_{(\cdot)}(\mathbf{x}) \equiv \frac{1}{\kappa_\infty(\mathbf{x})} \sum_\nu \frac{\kappa_\nu(\mathbf{x})}{\tau_\nu(\mathbf{x})} + \alpha_{(\cdot)}(\mathbf{x}) \quad (34)$$

and

$$\eta_\nu^{(\cdot)}(\mathbf{x}) \equiv \frac{\kappa_\nu(\mathbf{x})}{\tau_\nu(\mathbf{x})} - \kappa_\nu(\mathbf{x}) \alpha_{(\cdot)}(\mathbf{x}). \quad (35)$$

Also, analogous to the decomposition of the total pressure, the incident pressure is written as

$$p_i = p_{x,i} + p_{y,i} + p_{z,i}. \quad (36)$$

In principle, this decomposition of incident pressure introduces three more state variables than a corresponding solution for total pressure. In practice, however, only one additional state variable is required because the decomposed incident pressure is recombined at the beginning of each temporal iteration.

Finally, spatial and temporal staggering analogous to that described in Ref. 1 are used to improve computational accuracy. The three-dimensional extension of the spatially staggered two-dimensional grid defined in Ref. 1 is shown in Fig. 2. The spatial staggering and spatial derivatives are implemented using the so-called k -space operators¹ in a three-dimensional form that can be expressed as

FIG. 2. Spatially staggered grid.

$$\frac{\partial p(\mathbf{x}, t)}{\partial^{(c_0 \Delta t)^\pm} (\cdot)} \equiv \mathbb{F}^{-1} [j k_{(\cdot)} e^{\pm j k_{(\cdot)} \Delta(\cdot)/2} \text{sinc}(c_0 \Delta t k/2) \mathbb{F}[p(\mathbf{x}, t)]], \quad (37)$$

where $\mathbb{F}$ denotes a spatial Fourier transform, k is the magnitude of the wavenumber with components (k_x, k_y, k_z) , and the superscripted partial derivatives denote the dependence of the derivatives on step size and equilibrium wave speed. The k -space operators replace the spatial derivatives in Eqs. (32) and (33). As a result, the staggered variables are evaluated at points on three distinct spatial grids, one for each component of the pressure and velocity. The points are $\mathbf{x}_1 = \mathbf{x} + \mathbf{e}_x \Delta x/2$, $\mathbf{x}_2 = \mathbf{x} + \mathbf{e}_y \Delta y/2$, and $\mathbf{x}_3 = \mathbf{x} + \mathbf{e}_z \Delta z/2$ in which $\mathbf{e}_{(\cdot)}$ denotes a unit vector in the direction defined by the subscript. During a single time-step Δt , pressure is updated at time t , velocity is updated at times $t^\pm = t \pm \Delta t/2$, and the sign of the phase term in the k -space operator that implements spatial shifts is changed between velocity and pressure updates to maintain the same spatial grid. These k -space operators result in exact time stepping in a homogeneous medium and distinguish the parallel treatments in Refs. 1 and 22.

The resulting compact first-order spatially and temporally discrete equations for pressure and particle velocity of a scattered wave propagating in a medium that has relaxation absorption and is surrounded by a perfectly matched boundary layer for the total fields are

$$\begin{aligned} v_x(\mathbf{x}_1, t^+) &= e^{-\alpha_x(\mathbf{x}_1) \Delta t/2} \left\{ e^{-\alpha_x(\mathbf{x}_1) \Delta t/2} v_x(\mathbf{x}_1, t^-) \right. \\ &\quad \left. - \frac{\Delta t}{\rho(\mathbf{x}_1)} \frac{\partial p(\mathbf{x}, t)}{\partial^{(c_0 \Delta t)^+} x} - \left[\frac{\Delta t}{\rho(\mathbf{x}_1)} - \frac{\Delta t}{\rho_0} \right] \frac{\partial p_i(\mathbf{x}_1, t)}{\partial x} \right\}, \\ v_y(\mathbf{x}_2, t^+) &= e^{-\alpha_y(\mathbf{x}_2) \Delta t/2} \left\{ e^{-\alpha_y(\mathbf{x}_2) \Delta t/2} v_y(\mathbf{x}_2, t^-) \right. \\ &\quad \left. - \frac{\Delta t}{\rho(\mathbf{x}_2)} \frac{\partial p(\mathbf{x}, t)}{\partial^{(c_0 \Delta t)^+} y} - \left[\frac{\Delta t}{\rho(\mathbf{x}_2)} - \frac{\Delta t}{\rho_0} \right] \frac{\partial p_i(\mathbf{x}_2, t)}{\partial y} \right\}, \end{aligned}$$

$$v_z(\mathbf{x}_3, t^+) = e^{-\alpha_z(\mathbf{x}_3)\Delta t/2} \left\{ e^{-\alpha_z(\mathbf{x}_3)\Delta t/2} v_z(\mathbf{x}_3, t^-) - \frac{\Delta t}{\rho(\mathbf{x}_3)} \frac{\partial p(\mathbf{x}, t)}{\partial (c_0 \Delta t)^+} - \left[\frac{\Delta t}{\rho(\mathbf{x}_3)} - \frac{\Delta t}{\rho_0} \right] \frac{\partial p_i(\mathbf{x}_3, t)}{\partial z} \right\}, \quad (38)$$

$$\begin{aligned} p_x(\mathbf{x}, t + \Delta t) &= e^{-\mu_x(\mathbf{x})\Delta t} p_x(\mathbf{x}, t) - \frac{\Delta t}{\kappa_\infty(\mathbf{x})} e^{-\mu_x(\mathbf{x})\Delta t/2} \left[F_x(\mathbf{x}, t^+) + \frac{\partial v_x(\mathbf{x}_1, t^+)}{\partial (c_0 \Delta t)^+} - \sum_\nu \eta_\nu^x(\mathbf{x}) S_\nu^x(\mathbf{x}, t^+) \right], \\ p_y(\mathbf{x}, t + \Delta t) &= e^{-\mu_y(\mathbf{x})\Delta t} p_y(\mathbf{x}, t) - \frac{\Delta t}{\kappa_\infty(\mathbf{x})} e^{-\mu_y(\mathbf{x})\Delta t/2} \left[F_y(\mathbf{x}, t^+) + \frac{\partial v_y(\mathbf{x}_2, t^+)}{\partial (c_0 \Delta t)^+} - \sum_\nu \eta_\nu^y(\mathbf{x}) S_\nu^y(\mathbf{x}, t^+) \right], \\ p_z(\mathbf{x}, t + \Delta t) &= e^{-\mu_z(\mathbf{x})\Delta t} p_z(\mathbf{x}, t) - \frac{\Delta t}{\kappa_\infty(\mathbf{x})} e^{-\mu_z(\mathbf{x})\Delta t/2} \left[F_z(\mathbf{x}, t^+) + \frac{\partial v_z(\mathbf{x}_3, t^+)}{\partial (c_0 \Delta t)^+} - \sum_\nu \eta_\nu^z(\mathbf{x}) S_\nu^z(\mathbf{x}, t^+) \right], \end{aligned} \quad (39)$$

where

$$\begin{aligned} F_x(\mathbf{x}, t^+) &= [e^{\mu_x(\mathbf{x})\Delta t/2} \kappa_\infty(\mathbf{x}) - \kappa_0] \frac{\partial p_{x,i}(\mathbf{x}, t^+)}{\partial t} + \frac{\kappa_\infty(\mathbf{x})}{\Delta t} [e^{\mu_x(\mathbf{x})\Delta t/2} - e^{-\mu_x(\mathbf{x})\Delta t/2}] p_{x,i}(\mathbf{x}, t) - \sum_\nu \eta_\nu^x(\mathbf{x}) S_\nu^{x,i}(\mathbf{x}, t^+), \\ F_y(\mathbf{x}, t^+) &= [e^{\mu_y(\mathbf{x})\Delta t/2} \kappa_\infty(\mathbf{x}) - \kappa_0] \frac{\partial p_{y,i}(\mathbf{x}, t^+)}{\partial t} + \frac{\kappa_\infty(\mathbf{x})}{\Delta t} [e^{\mu_y(\mathbf{x})\Delta t/2} - e^{-\mu_y(\mathbf{x})\Delta t/2}] p_{y,i}(\mathbf{x}, t) - \sum_\nu \eta_\nu^y(\mathbf{x}) S_\nu^{y,i}(\mathbf{x}, t^+), \\ F_z(\mathbf{x}, t^+) &= [e^{\mu_z(\mathbf{x})\Delta t/2} \kappa_\infty(\mathbf{x}) - \kappa_0] \frac{\partial p_{z,i}(\mathbf{x}, t^+)}{\partial t} + \frac{\kappa_\infty(\mathbf{x})}{\Delta t} [e^{\mu_z(\mathbf{x})\Delta t/2} - e^{-\mu_z(\mathbf{x})\Delta t/2}] p_{z,i}(\mathbf{x}, t) - \sum_\nu \eta_\nu^z(\mathbf{x}) S_\nu^{z,i}(\mathbf{x}, t^+), \end{aligned} \quad (40)$$

$$\begin{aligned} S_\nu^x(\mathbf{x}, t^+) &= e^{-\Delta t/[2\tau_\nu(\mathbf{x})]} \left[e^{-\Delta t/[2\tau_\nu(\mathbf{x})]} S_\nu^x(\mathbf{x}, t^-) + \Delta t \frac{p_x(\mathbf{x}, t)}{\tau_\nu(\mathbf{x})} \right], \\ S_\nu^y(\mathbf{x}, t^+) &= e^{-\Delta t/[2\tau_\nu(\mathbf{x})]} \left[e^{-\Delta t/[2\tau_\nu(\mathbf{x})]} S_\nu^y(\mathbf{x}, t^-) + \Delta t \frac{p_y(\mathbf{x}, t)}{\tau_\nu(\mathbf{x})} \right], \end{aligned}$$

$$S_\nu^z(\mathbf{x}, t^+) = e^{-\Delta t/[2\tau_\nu(\mathbf{x})]} \left[e^{-\Delta t/[2\tau_\nu(\mathbf{x})]} S_\nu^z(\mathbf{x}, t^-) + \Delta t \frac{p_z(\mathbf{x}, t)}{\tau_\nu(\mathbf{x})} \right], \quad (41)$$

and

$$\begin{aligned} S_\nu^{x,i}(\mathbf{x}, t^+) &= e^{-\Delta t/[2\tau_\nu(\mathbf{x})]} \left[e^{-\Delta t/[2\tau_\nu(\mathbf{x})]} S_\nu^{x,i}(\mathbf{x}, t^-) + \Delta t \frac{p_{x,i}(\mathbf{x}, t)}{\tau_\nu(\mathbf{x})} \right], \\ S_\nu^{y,i}(\mathbf{x}, t^+) &= e^{-\Delta t/[2\tau_\nu(\mathbf{x})]} \left[e^{-\Delta t/[2\tau_\nu(\mathbf{x})]} S_\nu^{y,i}(\mathbf{x}, t^-) + \Delta t \frac{p_{y,i}(\mathbf{x}, t)}{\tau_\nu(\mathbf{x})} \right], \\ S_\nu^{z,i}(\mathbf{x}, t^+) &= e^{-\Delta t/[2\tau_\nu(\mathbf{x})]} \left[e^{-\Delta t/[2\tau_\nu(\mathbf{x})]} S_\nu^{z,i}(\mathbf{x}, t^-) + \Delta t \frac{p_{z,i}(\mathbf{x}, t)}{\tau_\nu(\mathbf{x})} \right]. \end{aligned} \quad (42)$$

III. NUMERICAL METHODS

The methods in this paper not only have the goal of performing accurate calculations of propagation over a large scale but also have the goal of utilizing efficiently a computing resource that is often shared among many users and on which access time is typically limited. Therefore, methods that allow optimal selection of calculation parameters are important. Examination of accuracy used media with simple geometries in which exact solutions to the forward problem were available for comparison. The k -space method was shown in two dimensions to be more accurate than other finite-difference time-domain methods.¹ Therefore, comparisons here were only made with exact methods. The codes used in this study for numerical solution of the forward problem in three dimensions were developed concurrently by the collaborating institutions with some interaction.^b The two codes, one that computed total pressure and one that computed scattered pressure, were shown to produce exactly the same scattering given an identical medium.

Because prospective applications of the k -space method are in the general area of biomedical ultrasound, accuracy and stability were evaluated with wideband plane-wave pulses characteristic of typical clinical and high-frequency biomedical ultrasound transducers. The clinical-frequency band is representative of bands used for diagnostic instrumentation and the high-frequency band is representative of bands used to study small animals. The clinical-frequency pulse was a Gaussian-windowed sinusoid with a 2.5-MHz center frequency and a -6 -dB bandwidth of 1.7 MHz. The high-frequency pulse was a Gaussian-windowed sinusoid with a 40-MHz center frequency and a -6 -dB bandwidth of 24 MHz.

A. Parameters of the PML

The PML allows the inhomogeneity to fill the entire computational domain by essentially eliminating reverberations but the matching region can extend computation time by increasing the size of the computational domain. Therefore, selection of PML parameters is important. In this study, the performance of the PML was evaluated in the high-frequency band as a function of the thickness of the PML and

the maximum absorption per grid point. A homogeneous medium was used with acoustic properties of water at body temperature ($c=1524$ m/s and $\rho=997$ kg/m³). The size of the spatial step was four points per minimum wavelength (PPW). This wavelength was computed at the upper -40 -dB frequency of the incident pulse and the time-step was chosen to result in a Courant–Friedrichs–Lewy (CFL) (Ref. 24) number of 0.5. Values of the CFL number reported here are based on the background sound speed while the time-step was determined by considering stability for a CFL number based on the maximum sound speed in the medium. The dispersionless PML attenuation parameters increased smoothly from the inner surface to the outer surface of the PML according to Eq. (27) in Ref. 1. The maximum absorption per grid point in the PML, $\alpha_{\max}$, was varied from 1 to 6 Np per grid point in increments of 0.5. The thickness of the PML was varied between 2 and 20 grid points in increments of 2. The pulse was propagated normally to the surface of the PML and the maximum reflection and transmission coefficients were determined as the PML thickness and absorption per PML point were varied.

B. Investigation of smoothing

Discontinuities in medium parameters, i.e., density, compressibility, and attenuation, produce, as previously noted, high spatial-frequency components in the medium spectrum and can cause aliasing when the first-order k -space operators given in Eq. (37) are evaluated. To reduce this effect, the medium parameters in some cases were lowpass filtered using a half-band filter.¹⁶ The benefit from smoothing was evaluated by comparing scattering of a clinical-frequency pulse from a small high-contrast 1-mm diameter sphere that was half-band filtered with scattering from a sphere that was not smoothed. This size sphere was chosen because the minimum wavelength in the studied clinical temporal-frequency band associated with the highest temporal frequency is 0.3 mm; hence, the sphere is small in the sense that it has a size of the order of (or smaller than) the incident wavelength. This property along with the stated spatial sampling criteria means that about 26 spatial samples span the diameter of the sphere when the sampling increment $\Delta x=0.075$ mm. If smoothing causes the radius of the sphere to grow by just two spatial samples, a 25% increase in the sphere volume results. Since scattering from fluid spheres in the long-wavelength limit is proportional to the sphere volume,²⁵ smoothing small spheres will be more likely to degrade agreement with an exact calculation. Therefore, an unfiltered medium at this scale may be more desirable, making the comparison of scattering from the smoothed medium with that of the unsmoothed medium at this scale a worst-case comparison; smoothing can be expected to perform better for larger scales. The parameters of the small sphere (i.e., sound speed, density, and attenuation slope) were chosen to implement a high-contrast object because small-scale inhomogeneities in human tissues are likely to be associated with microcalcifications.²⁶ To arrive at the parameters, the sound speed of a lossless sphere was increased from a background $c=1524$ m/s (water at body temperature) by increments of

300 m/s until an L^2 error (i.e., the spectral norm of the scaled difference between measured and exact scattering²) measured relative to an exact calculation exceeded 0.03. The high-contrast sound speed was set to the value that resulted in the excess L^2 error minus 150 m/s. Next, absorption was added to the sphere and the density increased from the background 997 kg/m³ in increments of 300 kg/m³ until the L^2 error doubled. The high-contrast density was equal to the density closest to the point at which this occurred.

For a scattering medium such as a single sphere or collection of spheres having an analytic Fourier transform, the filtering is exact. For arbitrary media constructed by segmentation of digitized images acquired using another modality (e.g., magnetic resonance imaging), the Fourier transform must be evaluated numerically. Accuracy in each case is controlled by the sampling rate. The sampling rate appropriate for an arbitrarily constructed medium was determined in this work by increasing the amount of oversampling for the sphere with parameters described in Sec. III D until the scattering from the oversampled and half-band filtered medium agreed with scattering from the exactly half-band filtered medium. Also, scattering from a lossless high-contrast sphere ($c_1=3540$ m/s, $\rho_1=1990$ kg/m³) in the same background with a volume equivalent to the volume of a single grid-point voxel, i.e., a so-called minimum-resolution sphere, was calculated for oversampling in multiples of 2, 4, 8, and 10.

C. Inclusion of absorption

The description of absorption caused by relaxation is for any finite number of relaxation processes. However, for each relaxation process, generalized compressibility-coefficient arrays for the medium require allocation of system memory. Thus, representation of attenuation with the minimum number of relaxation processes is desirable to model realistic absorption in tissue. Experience has shown that two processes satisfactorily model a linear dependence of absorption on frequency over the bands simulated. For clinical-frequency computations, relaxation times of 40 and 400 ns were used. For high-frequency computations, relaxation times of 2 and 20 ns were used. The long-duration relaxation time was chosen to be approximately the reciprocal of the lower -6 -dB frequency of the pulse and the short-duration relaxation time was chosen to be approximately five times the maximum frequency contained in the pulse. Compressibility coefficients for a lossy background were found by a least-square-error fit of Eqs. (42) and (43) in Ref. 7 for an attenuation coefficient slope with a linear dependence on temporal frequency. Fitting was performed over a frequency range of $0 < f < 5$ MHz for clinical pulses and $9.1 < f < 70.9$ MHz for high-frequency pulses. Evaluation of relaxation absorption was accomplished by propagating high-frequency Gaussian weighted pulses in the lossy medium with $c_0=1524$ m/s, $\rho_0=993$ kg/m³, and $\beta_0=0.5$ dB/(cm MHz) over a distance of the order of the maximum wavelength present in the pulse. Frequency-dependent attenuation derived from ratios of the pulse spectra and frequency-dependent phase velocities derived from

phase changes between the pulses at measurement points were compared with theoretical values from Ref. 7.

D. Accuracy and stability: L^2 vs CFL and PPW for a small sphere

A good compromise between CFL number and PPW was sought for large-scale calculations and comparisons with exact solutions. A low PPW is desirable to minimize spatial grid size because smaller grids required less storage and the three-dimensional FFTs can be evaluated more rapidly. Large values of CFL number are desirable because they require fewer temporal iterations to collect all non-negligible scattered waves. To determine the CFL and PPW values used for large-scale calculations, scattering from a test sphere was computed for selected values of PPW in the range 2–10 and CFL in the range 0.1–0.8. These ranges guarantee numerical stability and Nyquist temporal sampling. The sphere was 4 mm in diameter with sound speed $c_1=1460$ m/s, density $\rho_1=970$ kg/m³, and attenuation slope $\beta_1=0.5$ dB/(cm MHz) in a lossless background with sound speed $c_0=1509$ m/s and density $\rho_0=997$ kg/m³. The sphere was centered in the medium and excited by a clinical plane-wave pulse traveling in the $+e_x$ direction. Rotational symmetry about the axis defined by the incident wave vector permits the L^2 error of the scattering calculation to be found from the time-domain waveforms observed at points on a ring situated in the $z=0$ plane 2.5 mm from the center of the sphere. These time-domain waveforms were examined to find an L^2 error that produced agreement with an exact solution. The L^2 errors from the set of calculations along with sampling constraints and medium representation considerations were used to guide selection of PPW (Δx) and CFL (Δt) for large-scale comparisons.

E. Impact of inhomogeneity scale

Using a fixed PPW and CFL number determined with methods just described, the impact of increasing the scale of the calculations was studied by scattering clinical pulses from each of two sets of centered spheres in a lossless background with sound speed $c_0=1509$ m/s and density $\rho_0=997$ kg/m³. Both sets had radii ranging from 1 to 5 mm, which is the range of radii of spheres to be included in a large compound object discussed later. In one set, the spheres were lossless with sound speed $c_1=1570$ m/s and density $\rho_1=970$ kg/m³. In the other set, the spheres were lossy with $\beta_1=0.5$ dB/(cm MHz). The L^2 -error variation as a function of radius was tabulated.

Two more sets of spheres (0.75–4.55 mm) with “full contrast,” meaning lossy, i.e., $\beta_1=0.5$ dB/(cm MHz), with higher sound-speed contrast $c_1=1650$ m/s and density $\rho_1=970$ kg/m³ in a lossless background with sound speed $c_0=1509$ m/s and density $\rho_0=997$ kg/m³ were excited by a clinical pulse using PPW=4 and CFL $\approx$ 0.4. The L^2 error was calculated. Next, the same numerical experiment was performed at half the CFL number to demonstrate how accuracy was recovered at larger scales. These full contrast spheres were used to construct a compound scattering object described next.

TABLE I. Properties of background and spheres for large-scale compound medium.

Medium	Physical parameter		
	Sound speed (mm/ μ s)	Density (g/cc)	Absorption [dB/(cm MHz)]
Background	1.509	997	0.0
Large sphere	1.570	970	0.3
Smaller spheres	1.650	970	0.5

F. Large-scale compound object

A large-scale inhomogeneity with tissue-like properties^{27,28} was constructed by randomly positioning 12 non-overlapping spheres within a 24-mm diameter sphere. Calculation of scattering on such a scale required accurate wave propagation over nearly 300 wavelengths at the maximum temporal frequency contained in the pulse. The enclosed spheres all had the same acoustic parameters that differed from the acoustic parameters of the enclosing sphere. These acoustic parameters are reported in Table I. The radii and positions of the spheres are reported in Table II. Scattering was calculated from each centered unique sphere (eight different radii in all) as well as the compound sphere. The same data were collected from exact computations and the L^2 errors were tabulated. All scattering used four PPW and a CFL number approximately equal to 0.2 because all inhomogeneities were lossy and accurate large-scale calculation was the objective. For the compound medium, a comparison with an exact solution at a single temporal frequency ($f=2.5$ MHz) was made at 2943 points 14 mm from the center of the object by calculating the normalized root-mean-square error in the scattered-pressure magnitude.

IV. NUMERICAL RESULTS

A. PML parameters

The reflection and transmission coefficients for the PML parameter study are presented in Fig. 3 as functions of PML

TABLE II. Positions and radii of spheres forming the large-scale compound medium.

Radius (mm)	Position		
	x (mm)	y (mm)	z (mm)
12.00	0.00	0.00	0.00
4.55	−6.25	−2.50	−2.5768
3.90	−5.00	5.00	3.60
3.90	2.50	6.00	−3.60
2.50	7.00	1.00	2.00
2.50	2.50	−5.00	−2.00
2.50	7.00	6.00	1.50
2.50	−2.50	−9.00	−1.50
2.00	0.00	0.00	0.00
1.50	8.50	−3.50	0.00
1.25	3.75	−9.00	−1.00
1.25	−9.00	2.00	1.00
0.75	−2.00	10.00	0.00

FIG. 3. Left panel: Behavior of a PML as the maximum attenuation per cell is varied with the number of PML layers fixed at 10. Right panel: Behavior of a PML as the number of PML layers is varied and the maximum attenuation per PML cell is fixed at 3 Np.

thickness and maximum absorption per PML point. The right panel of Fig. 3 shows reflection and transmission coefficients vs the number of PML points when the maximum absorption per PML cell is 3 Np. The reflection and transmission coefficients both decrease as the thickness of the PML region increases but the marginal benefit of including more than 12 layers is small. Reflection and transmission coefficients are plotted as functions of maximum absorption per PML point in the left panel of the figure with PML thickness fixed at ten points. The transmission coefficient decreases monotonically as maximum absorption increases up to 6 Np per PML point but the rate of decrease is small once the maximum absorption exceeds 3.5 Np/point. In contrast, the reflection coefficient reaches a minimum at 2.5 Np per PML point. Similar behavior of the PML reflection coefficient as a function of maximum absorption per point was demonstrated in the original electromagnetic computations by Berenger.⁶ Figure 3 indicates that PML transmission and reflection coefficients less than -100-dB can be achieved simultaneously by setting the PML thickness to ten points and the maximum absorption to 3 Np per PML point. These PML parameters were employed in the simulations presented in this paper.

B. Medium smoothing

A worst-case comparison of scattering from a filtered and unfiltered high-contrast sphere with parameters $c_1=2874$ m/s, $\rho_1=2497$ kg/m³, and $\beta_1=0.5$ dB/(cm MHz) resulted in L^2 errors of 0.054 and 0.060, respectively. The scattered waveforms are shown in Fig. 4. Little difference exists between the waveforms for the two cases. This is consistent with the approximately equal errors except that the scattered waves in the unfiltered medium have a low-level background artifact typical for scattering from unfiltered media.

For a weak-scattering 2-mm radius sphere, oversampling a space-domain constructed medium by a factor of 2 using a minimum integer multiple for a half-band filter practically eliminates low-level artifacts such as the kind seen in Fig. 4. The L^2 error was insensitive to the amount of oversampling. For the high-contrast so-called minimum-resolution sphere,

oversampling improves agreement between scattering from the medium constructed exactly in the spatial-frequency domain with half-band filtering and the medium constructed by oversampling in the space domain followed by half-band filtering. The L^2 error is, however, lower for scattering computed from an unfiltered medium. This is attributed to increasing the effective volume of the scatterer.

In all cases, spatial filtering exactly or spatial filtering an oversampled domain by at least a factor of 2 practically eliminates artifact due to medium discontinuities. For small (grid size and smaller) inhomogeneities, spatial filtering broadens the spatial extent of the inhomogeneity and in-

FIG. 4. The scattered waves from a filtered (left) and unfiltered (right) 1-mm radius sphere. The horizontal axes in each panel span 360° in the $z=0$ plane with measurements taken at a distance of 3 mm from the sphere center and waveforms are shown on a bipolar logarithmic scale over a ± 60 dB range. The vertical axes span approximately 4.5 μ s.

FIG. 5. Temporal-frequency dependence of sound speed (left) and attenuation (right) in the high-frequency band.

creases error. For large-scale high-contrast inhomogeneities, such as a 2-mm sphere with the properties of a calcification, half-band filtering has the effect of slightly decreasing the overall L^2 error because some loss of accuracy for scattering in the forward direction is compensated by a greater improvement in backscatter accuracy. Backscattering from the filtered medium agrees better with an exact solution than the unfiltered medium. Another possible approach to numerical evaluation of smoothing is to compare errors in regions where the exact solution is quiescent. Artifact from unfiltered media increases errors in these regions and shows the need for filtering. Additionally, the energy level of the artifacts can be compared to energy levels expected from scattering by minimum-resolution (single-pixel) objects to determine if artifacts are larger than the actual scattering in the medium.

C. Absorption

The right panel of Fig. 5 illustrates the ability of the first-order k -space method with second-order relaxation absorption to model an approximately linear frequency dependence of attenuation with $\beta=0.5$ dB/(cm MHz) in the high-frequency band. In this figure, the theoretical second-order relaxation absorption model is compared to numerical results using CFL numbers of 0.5 and 0.25. The two-process model produces small deviations from a linear frequency dependence. With a CFL of 0.5, the k -space method produces attenuation within a few dB/cm of the theoretical curve at all frequencies in the -6 -dB bandwidth of the high-frequency pulse. Reducing the CFL number to 0.25 yields numerical attenuation very close to the theoretical curve. Fitting a linear dependence of attenuation on frequency produces a frequency dependence for sound speed shown in the left panel of Fig. 5. A comparison of numerical results to theoretical results shows good agreement for a CFL of 0.5 and excellent agreement for a CFL of 0.25. Analogous results are obtained at clinical frequencies.

D. Accuracy vs PPW and CFL for a small sphere

Scattering from a small sphere for a combination of grid sizes and time-steps yields the L^2 errors presented in Fig. 6. The results show that, as the grid spacing is decreased, smaller time-steps associated with lower CFL numbers are required to maintain accuracy. For these scattering experiments, the time-domain waveforms indicate that an L^2 error of approximately 0.02 or smaller is associated with excellent agreement between an exact solution and the k -space solution. Therefore, $L^2=0.02$ was used to guide selection of parameters for large-scale calculations.

A primary consideration for large-scale calculations is execution time. Execution-time results, presented in the Appendix, show that execution time is proportional to the number of time-steps and a power of the grid size. Therefore, reducing the grid size to the minimum acceptable value is desirable to reduce the computing system requirements. Nyquist sampling of a monochromatic wave at the highest temporal frequency of the clinical pulse would suggest Δx

FIG. 6. L^2 error resulting from scattering of clinical pulse waveforms by a 4-mm diameter sphere.

FIG. 7. Scattering from sets of spheres with progressively larger radii.

$=0.15$ mm (PPW=2) for large-scale calculations. However, because previously derived² stability criteria ignored density contrast that is included in the models used in the present study and no calculation for PPW=2 produced an L^2 error less than 0.02, the spatial sampling rate was doubled. An added benefit of doubling the spatial sampling rate is that discontinuities in the medium can be more accurately approximated and smaller inhomogeneities resolved. Given these considerations, a grid spacing of $\Delta x=0.075$ mm was selected for large-scale calculations of scattering of clinical pulses by spheres, and collections of spheres. With the selected grid spacing and given that $L^2 < 0.02$ produces accurate agreement with exact solutions, a time-step of $\Delta t = 0.02$ μs (CFL ≈ 0.55) is considered to be an upper bound for large-scale calculations.

E. Scaling results

Numerical scattering experiments with the two sets of spheres described in Sec. III E using prospective values for time-step and spatial grid size and using methods described in Sec. III D show that loss of accuracy can occur in both lossless and lossy media. The L^2 error as a function of scale size is plotted in Fig. 7 and shows that, as the sphere radius is varied from 1 to 5 mm, the accuracy of lossless scattering is weakly dependent on the scale size while lossy scattering shows a stronger dependence. As the size of the sphere is increased above 2 mm, accuracy decreases.

Also shown in Fig. 7 are L^2 errors in calculations of scattering from the inclusions in the large-scale compound object. Two different time-steps that result in CFL 0.4 and CFL 0.2 show that, at larger scales, discrepancies between the exact and k -space dispersion relationships enhance deviations from the exact solution. The L^2 errors associated with scattering at CFL 0.2 from the inclusions, the enclosing sphere, and the large-scale compound object are reported in Table III. Calculations for the large sphere and compound sphere were performed with a CFL number of 0.2 to maintain good agreement at large scales. Since error in the com-

TABLE III. L^2 errors of propagation calculations for each radius sphere in the large-scale compound medium.

Radius (mm)	L^2 error
0.75	0.007
1.25	0.005
1.50	0.005
2.00	0.006
2.50	0.009
3.90	0.021
4.55	0.018
12.00	0.021
12.00 ^a	0.025 ^b

^aCompound object.

^bSingle temporal frequency.

pound object is comparable to error observed in the inclusions, the error appears to be independent of the configuration of the scattering medium.

Snapshots of the scattered pressure in the $x, y, z=0$ planes at representative instants of time for the large-scale compound object are shown in Fig. 8. In this calculation, the incident wave (a clinical pulse) propagates in the $+e_x$ direction and is offset at $x=-16$ mm at time $t=0$. Therefore, scattering is shown in the $y, z=0$ planes as the incident wave penetrates the sphere but no scattering in the $x=0$ plane occurs until the incident wave arrives at the center of the sphere. The entire video sequence (MPEG-2) is available at the webpage found in Ref. 29. The time histories allow appreciation of scattering into a plane from out-of-plane inhomogeneities.

V. DISCUSSION

A. Solutions for total pressure and scattered pressure

When the first-order wave equations are written in terms of the total pressure, the medium can be passively excited by specifying a propagating wave in the medium at time $t=0$. For this case, the medium must contain the incident wave as well as the inhomogeneities. The total-pressure forward-problem solution is, therefore, a natural choice for pulse-echo applications in which the initial position of the wave corresponds to receiver positions that must also be within the medium to sample spatially limited transmit and receive apertures. The total-pressure forward-problem solution can also be used to excite incident waves of infinite spatial extent. Such waves are created with wave vectors parallel to a Cartesian axis because reflections at the PML are produced for all incident directions that are not collinear with one of the Cartesian coordinate system axes. The angle of incidence can, however, be adjusted by rotating the inhomogeneity. Single-frequency and pulsatile incident plane waves can be implemented by specifying the values of pressure and velocity on the boundary of the grid. Such an excitation requires a time-domain calculation of the incident wave at every boundary point and values of the incident field are specified on the boundary at every time-step of the calculation. The incident pressure for the scattered-pressure form of the first-

FIG. 8. Scattering from a 24-mm diameter compound sphere with properties similar to those of human breast in a background of water is shown in three orthogonal planes. At the top are time frames recorded at $x=0$. In the middle are time frames recorded at $y=0$. At the bottom are time frames recorded at $z=0$. The six frames are recorded at times 4.8, 9.6, 10.4, 11.2, 12.4, and 17.2 μs . The incident illumination is a plane-wave pulse offset at $x=-16$ mm at time $t=0$ with center frequency of 2.5 MHz and a Gaussian envelope with a -3 dB bandwidth of 1.7 MHz. The medium sound speed is overlaid on the pressure field with brighter areas corresponding to higher sound speed.

order equations is, in contrast, always active. This allows implementation of single-frequency and pulsatile plane waves from arbitrary directions. The incident wave is, however, needed not just on the boundary but also at each grid point. This added burden is not significant for simple incident waves but can add computational cost and storage for more complex incident waves such as focused beams.

B. Scaling and accuracy

Figure 7 suggests that large-scale media degrade accuracy. Why this occurs can be understood with the help of Fig. 5 that compares the k -space dispersion relationships with the exact theoretical values. For theoretical calculations, the dispersion relationships are exact. For k -space method calculations, the actual dispersion relationships present in the k -space results can, however, deviate from the exact theoretical values. The agreement between k -space and exact dispersion relationships is best for lower CFL numbers. As the scale of the inhomogeneity increases, the discrepancy between k -space and exact calculations increases because scattered waves accumulate. For example, absorption causes spectral components of a pulse to suffer amplitude reductions that accumulate as the propagation distance in the lossy medium increases. In general, calculation of scattering in real tissues³⁰ accurately using the k -space method requires close agreement between the dispersion relationships.

A more rigorous method for selection of relaxation times is to search for relaxation times of a specific order that minimize the mean square error between a linear frequency dependence of attenuation and the resulting theoretical curve. If this is done for the clinical-frequency range, relaxation times

of 29 and 258 ns result. Although these are not the values chosen here, the comparisons with exact methods remain valid because exact codes used theoretical values based on the same relaxation times used in the k -space code. Two relaxation processes are justified (as opposed to three or more) because the rms error measured relative to the attenuation at 2.5 MHz (clinical center frequency) is only 1.3%. Inclusion of a third process only reduces the rms error measured relative to the attenuation at 2.5 MHz to 0.8%.

VI. CONCLUSIONS

The three-dimensional extension of the k -space code calculates scattering with an accuracy that is similar to that achieved with the two-dimensional code. Investigations with the two-dimensional code did not, however, include comparisons with exact scattering for lossy inhomogeneities. Therefore, because the k -space method is exact for lossless, homogeneous media, the departure of k -space calculated scattering from exact solutions as scale size increases was not previously observed. This behavior underscores the importance of using smaller CFL numbers to maintain accuracy for large-scale calculations that include loss when comparisons are made with results from exact methods. Similarly, if accurate comparisons are to be made with realistic models of tissue, then the distribution of k -space relaxation times must be chosen to model closely the dispersion relationships of the tissue to maximize agreement. The results of this paper show that calculation of propagation of ultrasound over hundreds of wavelengths is made feasible by the three-dimensional implementation of the k -space method on a cluster of computers such as those described. Increasing availability of suit-

able computing platforms will allow investigation, through simulation, of *b*-mode imaging and other applications such as ultrasound beam aberration caused by inhomogeneities in tissue. Furthermore, inverse problem solutions require a solution to the forward problem and the described method is a realistic approach to solving the forward problem for arbitrary configurations of tissue.

ACKNOWLEDGMENTS

David P. Duncan is thanked for discussions and code development. Jeffrey P. Astheimer and Jing Jin are thanked for participation in reviews throughout the duration of this work. Andrew J. Hesford is thanked for discussions during the final stages of the reported studies. Initial development of code used resources at Research Computing and the Department of Electrical Engineering, Multi-agent Bio-Robotics Lab, at the Rochester Institute of Technology. Gurcharan S. Khanna and Ferat Sahin are thanked for granting access to these resources. This research also used the SHARCNET facilities. Ge Baolai is thanked for technical assistance in the use of SHARCNET. Subsequent code development and computations used resources of the National Energy Research Scientific Computing Center (NERSC), which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. Consulting staff at NERSC are thanked for help compiling benchmark test programs on the Cray XT4 (Franklin). M.I.D. is a Scholar of the Canadian Institutes of Health Research/University of Western Ontario Strategic Training Initiative in Cancer Research and Technology Transfer and holds a NSERC PGS-D scholarship. This research was funded in part by NIH Grant Nos. HL 50855, CA 74050, and EB 00280, Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant Nos. 261323-03 and 261323-07, the Canadian Institutes of Health Research/University of Western Ontario Strategic Training Initiative in Cancer Research and Technology Transfer, and the University of Rochester Diagnostic Ultrasound Research Laboratory Industrial Associates.

APPENDIX: PARALLEL PERFORMANCE OF FFTW

Calculations were performed on two different clusters of computers. Performance results were compiled for both systems for comparison. The three-dimensional distributed FFT was the computational bottleneck so the execution time of a single three-dimensional FFT was analyzed on both systems as functions of the size of the computational grid and number of nodes using test programs that are distributed with the FFTW software. The grid size was varied from 32 to 2048 points in multiples of 2 simultaneously along each spatial dimension. The number of nodes was varied from 2 to 2048 in multiples of 2. The mean execution time for ten single-transform trials for both forward and backward transforms was noted for each combination of cluster, grid size, and number of nodes. The first cluster, Franklin (<http://www.nersc.gov/nusers/systems/franklin>), a Cray XT4 system at Lawrence Livermore National Laboratory, used nodes with 2.6 GHz dual-core AMD Opteron processors, 4 Gbytes of memory, and SeaStar2 (Refs. 31 and 32) interconnects

FIG. 9. (Color online) Execution times for the three-dimensional FFT on the Franklin and SHARCNET clusters.

running a Linux operating system. The second cluster, SHARCNET (<http://www.sharcnet.ca>), the Shared Hierarchical Academic Research Computing Network spread among several institutions in Ontario, used nodes with Opteron single-core 2.4-GHz processor, 32 Gbyte memory, and Quadrics Elan4 interconnects running an HP XC 3.1 Linux-based system. Since Franklin has many more compute nodes than SHARCNET, performance evaluations for some permutations in the set of possible grid sizes and nodes were not compiled on SHARCNET. The set of permutations of grid sizes and nodes selected ensured that adequate memory existed on each compute node to allocate arrays and minimize the probability of a node using a disk memory cache during the calculation.

A comparison of the execution time for a single distributed three-dimensional FFT on the two different computer clusters used to produce results in this paper is shown in Fig. 9. Runtimes on both clusters have scaling similarly with grid size for a fixed node count. The results show SHARCNET as slightly better scaling. However, for a fixed number of nodes, Franklin always performs the transform in less than half the time of SHARCNET. In all cases, for a fixed grid size, SHARCNET performs better by using more nodes. For Franklin, this is not the case. For some grid sizes, increasing the node count lengthens the execution time.

The difference in the slope of the execution-time lines between Franklin and SHARCNET for a fixed number of processors (two, for example) merits comment. The apparent lower slope of execution time on SHARCNET is attributed to node differences such as per-core memory. Faster execution on Franklin than on SHARCNET for any combination of nodes and grid size is attributed to Franklin's SeaStar2 high-performance, low-latency interconnects being available because the distributed three-dimensional FFT requires an all-to-all communication prior to performing the last set of two-dimensional transforms for the slab on each node. Execution time on Franklin can be different than on SHARCNET because performing a transform of fixed size with

fewer nodes is sometimes faster. The grid size at which this behavior occurs is proportional to the number of nodes performing the transform.

For a fixed number of processors, the time required to execute a single three-dimensional FFT scales as $O(N^3 \log N)$ for N^3 grid points. The benefit of using a cluster of computers results from increasing the number of nodes used to calculate the three-dimensional FFT as the grid size is increased. If the number of nodes is chosen to be equal to one dimension of a cubic medium, then the scaling exponent is reduced from 3 to 1.7. This benefit along with capabilities to calculate scattering in a large-scale medium that would otherwise be intractable because of the aggregate memory required for storage is an advantage of the parallel k -space method.

^aIn Eqs. (17) and (18), and in following analysis, the notation uses the quantum mechanics convention in which a symbol represents a conceptual object and the argument indicates the domain in which an object is evaluated. Thus, $p(\mathbf{x}, t)$ is the scattered pressure in the space domain as a function of time t and $p(\mathbf{x}, \omega)$ is the temporal Fourier transform as a function of temporal-frequency ω .

^bBoth institutions benefited from a previously developed two-dimensional algorithm from the University of Rochester and some initial exchanges concerning architecture but otherwise three-dimensional extensions of the code proceeded independently.

¹M. Tabei, T. D. Mast, and R. C. Waag, "A k -space method for coupled first-order acoustic propagation equations," *J. Acoust. Soc. Am.* **111**, 53–63 (2002).

²T. D. Mast, L. P. Souriau, D.-L. D. Liu, M. Tabei, A. I. Nachman, and R. C. Waag, "A k -space method for large-scale models of wave propagation in tissues," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control* **48**, 341–354 (2001).

³B. Comani-Tabrizi, " k -space formulation of the absorptive full fluid elastic scalar wave equation in the time domain," *J. Acoust. Soc. Am.* **79**, 901–905 (1986).

⁴S. Finette, "Computational methods for simulating ultrasonic scattering in soft tissue," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control* **34**, 283–292 (1987).

⁵S. Pourjavid and O. J. Tretiak, "Numerical solution of the direct scattering problem through the transformed acoustical wave equation," *J. Acoust. Soc. Am.* **91**, 639–645 (1992).

⁶J.-P. Berenger, "A perfectly matched layer for the absorption of electromagnetic waves," *J. Comput. Phys.* **114**, 185–200 (1994).

⁷A. I. Nachman, J. F. Smith, and R. C. Waag, "An equation for acoustic propagation in inhomogeneous media with relaxation losses," *J. Acoust. Soc. Am.* **88**, 1584–1595 (1990).

⁸N. N. Bojarski, "The k -space formulation of the scattering problem in the time domain," *J. Acoust. Soc. Am.* **72**, 570–584 (1982).

⁹N. N. Bojarski, "The k -space formulation of the scattering problem in the time domain: An improved single propagator formulation," *J. Acoust. Soc. Am.* **77**, 826–831 (1985).

¹⁰M. I. Daoud and J. C. Lacefield, "Distributed three-dimensional simulation of B-mode ultrasound imaging using a first-order k -space method," *Phys. Med. Biol.* (accepted for publication).

¹¹T. L. Sterling, J. Salmon, D. J. Becker, and D. F. Savarese, *How to Build a Beowulf* (MIT, Cambridge, MA, 1999).

¹²M. Frigo and S. G. Johnson, "The design and implementation of FFTW3," *Proc. IEEE* **93**, 216–231 (2005).

¹³O. M. Al-Bataineh, T. D. Mast, E. Park, V. W. Sparrow, R. M. Keolian, and N. B. Smith, "Utilization of the k -space method in the design of a ferroelectric hyperthermia phased array," *Ferroelectrics* **331**, 103–120 (2006).

¹⁴F. Jensen, W. Kuperman, M. Porter, and H. Schmidt, *Computational Ocean Acoustics* (Springer, New York, 2000).

¹⁵B. T. Cox and P. C. Beard, "Fast calculation of pulsed photoacoustic fields in fluids using k -space methods," *J. Acoust. Soc. Am.* **117**, 3616–3627 (2005).

¹⁶R. E. Crochiere and L. R. Rabiner, *Multirate Digital Signal Processing* (Prentice-Hall, Englewood Cliffs, NJ, 1983), pp. 155–157.

¹⁷P. M. Morse and K. U. Ingard, *Theoretical Acoustics* (McGraw-Hill, New York, 1968), Chap. 6.

¹⁸A. D. Pierce, *Acoustics: An Introduction to Its Physical Principles and Applications*, 2nd ed. (Acoustical Society of America, Woodbury, NY, 1989), Chap. 1.

¹⁹J. Meixner, "Absorption and dispersion of sound in gases with chemically reacting and excitable components," *Ann. Phys.* **5** (43), 470–487 (1943).

²⁰S. R. de Groot and P. Mazur, *Non-Equilibrium Thermodynamics* (North-Holland, Amsterdam, 1962), Chap. 2.

²¹A. D. Pierce, *Acoustics: An Introduction to Its Physical Principles and Applications*, 2nd ed. (Acoustical Society of America, Woodbury, NY, 1989), Chap. 10.

²²X. Yuan, D. Borup, J. W. Wiskin, M. Berggren, and S. A. Johnson, "Simulation of acoustic wave propagation in dispersive media with relaxation losses by using FDTD method with PML absorbing boundary condition," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control* **46**, 14–23 (1999).

²³X. 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. Ultrason. Ferroelectr. Freq. Control* **44**, 816–822 (1997).

²⁴E. Turkel, "On the practical use of high-order methods for hyperbolic systems," *J. Comput. Phys.* **35**, 319–340 (1980).

²⁵P. M. Morse and K. U. Ingard, *Theoretical Acoustics* (McGraw-Hill, New York, 1968), Chap. 8.

²⁶A. Thomas Stavros, *Breast Ultrasound* (Lippincott, Williams, and Wilkins, Philadelphia, PA, 2004), Chap. 12.

²⁷F. T. D'Astous and F. S. Foster, "Frequency dependence of ultrasound attenuation and backscatter in breast tissue," *Ultrasound Med. Biol.* **12**, 795–808 (1986).

²⁸Weiward, W., Heinig, A., Goetz, L., Hartmann, H., Lampe, D., Buchmann, J., Millner, R., Spiekman, R. P., and Heywang-Köbrunner, S. H. "Direct measurement of sound velocity in various specimens of breast tissue," *Invest. Radiol.* **35**, 721–726 (2000).

²⁹J. C. Tillett, "Scattering in three dimensions," <http://www.ece.rochester.edu/projects/ultrasounds/3dScattering.mpg>, MPEG-2 Video (Last viewed 5/5/2009).

³⁰A. Francis, *Duck, Physical Properties of Tissue* (Academic, San Diego, CA, 1990), Chap. 4.

³¹R. Alverson, "Red storm," an invited talk, Hot Chips, 15 August (2003); http://www.hotchips.org/archives/hc15/2_Mon/1.cray.pdf, Cray Red Storm system overview (Last viewed 12/31/2008).

³²R. Brightwell, K. T. Pedretti, K. D. Underwood, and T. Hudson, "SeaStar interconnect: Balanced bandwidth for scalable performance," *IEEE MICRO* **26**, 41–57 (2006).