Comparison of temporal and spectral scattering methods using acoustically large breast models derived from magnetic resonance images

Andrew J. Hesford,^{a)} Jason C. Tillett, and Jeffrey P. Astheimer

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

Robert C. Waag

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

(Received 17 June 2013; revised 31 January 2014; accepted 24 June 2014)

Accurate and efficient modeling of ultrasound propagation through realistic tissue models is important to many aspects of clinical ultrasound imaging. Simplified problems with known solutions are often used to study and validate numerical methods. Greater confidence in a time-domain k -space method and a frequency-domain fast multipole method is established in this paper by analyzing results for realistic models of the human breast. Models of breast tissue were produced by segmenting magnetic resonance images of *ex vivo* specimens into seven distinct tissue types. After confirming with histologic analysis by pathologists that the model structures mimicked *in vivo* breast, the tissue types were mapped to variations in sound speed and acoustic absorption. Calculations of acoustic scattering by the resulting model were performed on massively parallel supercomputer clusters using parallel implementations of the k -space method and the fast multipole method. The efficient use of these resources was confirmed by parallel efficiency and scalability studies using large-scale, realistic tissue models. Comparisons between the temporal and spectral results were performed in representative planes by Fourier transforming the temporal results. An RMS field error less than 3% throughout the model volume confirms the accuracy of the methods for modeling ultrasound propagation through human breast. © 2014 Acoustical Society of America. [<http://dx.doi.org/10.1121/1.4887461>]

PACS number(s): 43.58.Ta, 43.20.Fn, 43.80.Qf, 43.80.Ev [VWS]

Pages: 682–692

I. INTRODUCTION

A common and fundamental component of imaging research is the computation of acoustic scattering from model human tissue volumes that will be used to develop and test imaging algorithms as well as to study aberration of acoustic beams.^{1–4} Aberration studies have been and continue to be used to define resolution limits of clinical ultrasound imaging systems and to identify critical tissue characterization parameters derived from observed effects on ultrasound wave propagation through tissue. The outcomes of these studies will facilitate the clinical implementation of adaptive ultrasound focusing techniques that compensate for complicated aberration produced by tissue. Versions of the aberration correction algorithms have been verified in experiments to produce essentially diffraction-limited focusing.^{5,6} Quantification of resolution limits in the presence of aberration and with the use of aberration correction will lead to development of improved focusing techniques. These techniques could also be used to tailor transducer apertures to obtain the best balance between resolution and cost for commercial systems that image *in vivo* human tissues, resulting in improved diagnostic capabilities.

Another important application for forward solutions is in linearized iterative imaging algorithms such as the distorted Born iterative method^{7–10} and the mathematically equivalent eigenfunction method.^{11–13} These techniques require repeated evaluation of forward scattering for determining successive updates to estimates of the medium. Improved methods that are applicable in clinical settings require realistic scattering data to provide realistic evaluations of performance.

This paper seeks to establish the efficiency and validity of the well-known fast multipole method (FMM)^{14–16} and the k -space method^{17–19} for modeling ultrasound scattering by realistic models of human breast tissue. The FMM computes scattering from arbitrary media using integral equations at a fixed temporal frequency. With a computational complexity that is $O(N)$ for a medium described by N scattering elements, the FMM has already been employed to produce data for scattering problems of extremely large size. The k -space method computes scattering from arbitrary media using difference equations and a temporal field representation to compute solutions over a finite bandwidth. While the $O(N \log N)$ computational scaling of the k -space method is slightly worse than that of the FMM, the scaling constant is smaller so the method is still very efficient for the problems of interest. Furthermore, the availability of multiple-frequency solutions after a single run is a strong advantage of the k -space method. For this study, models of

^{a)}Author to whom correspondence should be addressed. Electronic mail: andrew.hesford@rochester.edu

human breast were segmented into seven tissue types using high-resolution, three-dimensional magnetic resonance images of *ex vivo* specimens. Previously published measurements of tissue acoustic parameters were used as a basis for a mapping that converted the segmented tissue types into values of sound speed and acoustic absorption for each voxel. The k -space method and the FMM were then used to simulate ultrasound scattering by the model to facilitate comparisons of the methods. Additionally, the large-scale, realistic breast model was used in a unique study of the parallel performance of the FMM for problems of clinical relevance. Similar studies of the k -space method have been previously described.^{20,21}

No known prior attempts have been made to validate either method using models of comparable acoustic size or complexity. Both of these aspects of the computed models present unique challenges that require novel formulations and implementations of acoustic scattering solvers. The FMM^{22,23} and the k -space method^{21,24,25} have each been substantially enhanced to improve efficiency and also to extend the scope of solvable problems characterized by segmented magnetic resonance images. Combinations of these prior improvements made possible the case studies presented here of previously infeasible FMM and k -space solutions to large-scale, realistic scattering problems. Because known scattering solutions for the computed models are unknown, the validity of the numerical solutions is established through close agreement of the two independent methods.

II. ACOUSTIC SCATTERING THEORY

Ultrasound waves within biological tissue are governed in three dimensions by the first-order acoustic equations²¹

$$\rho(\mathbf{r})\partial_t \mathbf{u}(\mathbf{r}, t) = -\nabla p(\mathbf{r}, t), \quad (1a)$$

$$\int_0^t d\tau \kappa(\mathbf{r}, t - \tau) \partial_\tau p(\mathbf{r}, \tau) = -\nabla \cdot \mathbf{u}(\mathbf{r}, t), \quad (1b)$$

where p denotes pressure, $\mathbf{u}$ denotes particle velocity and κ denotes compressibility. Equation (1a) is a form of Newton's second law and Eq. (1b) expresses the resistance of the medium to compression. The time variations of compressibility κ are causal at each spatial location and, consequently, the convolution in Eq. (1b) produces delayed responses to changes in pressure that are spread out in time.

A wave equation for the pressure variations is obtained from Eqs. (1) by taking the time derivative of Eq. (1b) and substituting $-\rho^{-1}\nabla p$ for $\partial_t \mathbf{u}(\mathbf{r}, t)$. This gives

$$\nabla \cdot [\rho^{-1}(\mathbf{r})\nabla p(\mathbf{r}, t)] - \kappa(\mathbf{r}, t) * \partial_{tt} p(\mathbf{r}, t) = 0 \quad (2)$$

in which the operator $*$ denotes temporal convolution. The sound speed c and the absorption coefficient γ are frequency dependent and do not explicitly appear in this temporal wave equation. However, taking the temporal Fourier transform of Eq. (2) and multiplying by ρ yields the monochromatic equation

$$\rho(\mathbf{r})\nabla \cdot [\rho^{-1}(\mathbf{r})\nabla p(\mathbf{r}, \omega)] + k^2(\mathbf{r}, \omega)p(\mathbf{r}, \omega) = 0, \quad (3)$$

where the dispersion relation

$$k(\mathbf{r}, \omega) = \omega \sqrt{\rho(\mathbf{r})\kappa(\mathbf{r}, \omega)} = \frac{\omega}{c(\mathbf{r}, \omega)} + i\gamma(\mathbf{r}, \omega) \quad (4)$$

relates the wave number k , sound speed c and absorption γ to the density and compressibility.

III. NUMERICAL METHODS

Temporal solutions of Eq. (2) are obtained using the k -space method described in Ref. 21, and monochromatic solutions of Eq. (3) are obtained efficiently using the fast multipole method.^{22,23} There are advantages and disadvantages to both methods. The k -space method produces a complete temporal solution that can be Fourier transformed to obtain monochromatic solutions at any spatial frequency. However, if only one or several monochromatic solutions are required, then the complete solution provided by the k -space method is more costly than computing separate FMM solutions at each of the selected frequencies.

A. The fast multipole method

Monochromatic scattering solutions are obtained by introducing a source term S into the wave equation (3) to obtain

$$\rho(\mathbf{r})\nabla \cdot [\rho^{-1}(\mathbf{r})\nabla p(\mathbf{r}, \omega)] + k^2(\mathbf{r}, \omega)p(\mathbf{r}, \omega) = -S(\mathbf{r}, \omega). \quad (5)$$

For the models described in this paper, density variations were observed to correspond to variations of less than 6% (RMS) in the acoustic pressure fields within the models. Comparisons of results from the frequency-domain and time-domain solutions were not significantly altered by removing density variations from the model. Therefore, to simplify more detailed analyses, density was assumed to be constant in all models. As a result, Eq. (5) reduces to the Helmholtz equation

$$[\nabla^2 + k_0^2]p(\mathbf{r}) = -S(\mathbf{r}) - O(\mathbf{r})p(\mathbf{r}), \quad (6)$$

where the wave number k_0 corresponds to the sound speed c_0 and absorption coefficient γ_0 of a homogeneous, ambient fluid in which the tissue sample is immersed, and O is called the object contrast, or scattering potential, given by

$$O(\mathbf{r}) = k^2(\mathbf{r}) - k_0^2. \quad (7)$$

The frequency dependence of p , S , k_0 , k , and O has been suppressed for brevity.

The integral form of Eq. (6) is

$$p(\mathbf{r}) - \int_D d\mathbf{r}' g_0(\mathbf{r}, \mathbf{r}') O(\mathbf{r}') p(\mathbf{r}') = p^i(\mathbf{r}) \quad (8)$$

in which g_0 is the free-space homogeneous Green's function and the incident field $p^i(\mathbf{r})$ satisfies

$$[\nabla^2 + k_0^2]p^i(\mathbf{r}) = -S(\mathbf{r}). \quad (9)$$

Numerical solutions of Eq. (8) proceed by subdividing the domain D into a collection $\{c_j; 0 \leq j < N\}$ of N disjoint, cubic cells. These cells coincide with the cells of segmented tissue models obtained from magnetic resonance images. Associated with each cell c_j are a constant contrast value O_j and a constant field value p_j such that

$$O(\mathbf{r}) = \sum_{j=0}^{N-1} O_j \chi_j(\mathbf{r}), \quad p(\mathbf{r}) = \sum_{j=0}^{N-1} p_j \chi_j(\mathbf{r}), \quad (10)$$

where χ_j here represents the characteristic, or pulse basis, function of cell c_j . Using the expansions (10) in Eq. (8) results in the matrix equation

$$p - G\tilde{p} = p^i \quad (11)$$

in which p and $\tilde{p}$ denote column vectors $[\Delta_j p_j]$ and $[O_j p_j]$, respectively. The column vector $p^i = [p_j^i]$ contains entries

$$p_j^i = \int_D d\mathbf{r} \chi_j(\mathbf{r}) p^i(\mathbf{r}) \quad (12)$$

and the Green's matrix $G = [G_{ij}]$ contains entries

$$G_{ij} = \int_D d\mathbf{r} \chi_i(\mathbf{r}) \int_D d\mathbf{r}' g_0(\mathbf{r}, \mathbf{r}') \chi_j(\mathbf{r}'). \quad (13)$$

These discrete representations correspond to a Galerkin testing scheme that ensures the matrix G is symmetric.

The FMM employs a recursive subdivision of the scattering domain into a hierarchy of successively smaller groups of elements among which interactions are evaluated. The discrete scattering represented by Eq. (11) is split into near- and far-field contributions to obtain

$$p - [G_F + G_N]\tilde{p} = p^i. \quad (14)$$

The operator G_F represents interactions between well-separated groups that may be computed using the FMM. The definition of "well-separated" depends on the desired computational accuracy, but here it refers to groups within the same level of the hierarchy that are not immediately adjacent. The key to efficiency in the FMM is representing interactions among groups of scattering elements at the coarsest possible level of the hierarchy. For groups that are not well separated in one level, well-separated subdivisions of the groups use the FMM to model interactions at the next finer level. In the finest level of the hierarchy, interactions among groups that are not well separated must be computed directly using Eq. (13); such direct interactions are represented by the operator G_N . A detailed description of the FMM and a discussion of the concept of well-separated groups appears in Ref. 16.

Because the FMM provides only a mechanism for representing the operator G_F in Eq. (14), iterative solution methods such as the generalized minimal residual method (GMRES)²⁶ or the stabilized biconjugate gradient method (BiCG-STAB)²⁷ are employed in concert with the FMM to solve Eq. (11) efficiently without the need to explicitly construct the Green's matrix (13). This pairing allows each

iteration to be computed in $O(N)$ time and with $O(N)$ storage. GMRES was used in the studies described below.

B. The k -space method

The k -space method is a time-step algorithm that advances the spatial pressure field in time according to a discretized representation of Eq. (2). The pressure, along with a set of auxiliary parameters, is computed at each time instant throughout the volume of interest from values of the pressure and the auxiliary parameters at previous times. Since the volume of interest is finite, precautions must be taken to prevent reflection from the boundary of the volume. This is accomplished by appending a perfectly matched boundary layer (PML) to the volume that absorbs the wave components that are normal to the boundary.^{24,28} Implementation of the PML requires a more general formulation of Eqs. (1) in which the total pressure is decomposed into a sum $p = p_x + p_y + p_z$ of negative stresses in the x , y and z directions. In that case, Eq. (1b) splits into the three stress-strain relationships

$$\kappa(\mathbf{r}, t) * \partial_t p_\eta(\mathbf{r}, t) = -\partial_\eta u_\eta(\mathbf{r}, t), \quad (15)$$

in which $\eta = x$, $\eta = y$ or $\eta = z$.

The total pressure p obtained by solving the system of differential equations given by Eqs. (1a) and (15) is the same as the pressure obtained by solving the original system of equations given by Eqs. (1). However, directional attenuation parameters α_x , α_y , and α_z in the boundary layer can be appended to the more general equations to yield²¹

$$\rho(\mathbf{r})[\partial_t u_\eta(\mathbf{r}, t) + \alpha_\eta(\mathbf{r})u_\eta(\mathbf{r}, t)] = -\partial_\eta p(\mathbf{r}, t), \quad (16a)$$

$$\kappa(\mathbf{r}, t) * [\partial_t p_\eta(\mathbf{r}, t) + \alpha_\eta(\mathbf{r})p_\eta(\mathbf{r}, t)] = -\partial_\eta u_\eta(\mathbf{r}, t). \quad (16b)$$

Because the parameters α_η vanish outside of the PML, this system is equivalent to Eq. (2) in the volume of interest.

The compressibility κ in Eqs. (2), (4), and (15) is given by an explicit relaxation expansion

$$\kappa(\mathbf{r}, t) = \kappa_\infty(\mathbf{r})\delta(t) + \sum_\nu \frac{\kappa_\nu(\mathbf{r})}{\tau_\nu(\mathbf{r})} e^{-t/\tau_\nu(\mathbf{r})} H(t) \quad (17)$$

for the compressibility function also given in Ref. 21, where κ_∞ and κ_ν are compressibility amplitudes, τ_ν are relaxation times, and $H(t)$ is the Heaviside function. Equation (17) provides a mechanism for incorporating observed variation in absorption with frequency and is derived from a thermodynamic model which characterizes the tissue as a mixture of species with different chemical affinities. While tissues likely exhibit many relaxation processes for energy absorption, a small number may be used provided that the processes can model observed linear frequency dependence of absorption in tissue. An earlier study²¹ has shown that two processes are adequate for modeling linear frequency dependence for absorption so two processes are assumed for calculations described here. However, previous work^{1,21} assumed that the generalized compressibility amplitudes κ_ν and κ_∞ , which must be calculated for each triplet of density, attenuation slope β (measured in dB/cm/MHz), and sound

speed in the medium, could be derived by scaling reference values. Scaling was used because for each unique triplet, a computationally intensive, least-squares method would otherwise be used to calculate the coefficients that best fit a target attenuation slope.

The assignment of acoustic parameters to a segmented magnetic resonance image of the breast uses reported measurements that are typically performed with broadband acoustic transducers. The values for parameters in particular tissue types should be interpreted as those corresponding to the center frequency of the measurement. For single-frequency calculations at this center frequency, no disparity exists between reported tissue properties and assigned model parameters. However, the compressibility model (17) used in the k -space method introduces dispersion when absorption is present. Values for the coefficients κ_∞ and κ_ν that yield the desired attenuation slope over the band of the pulse, and that match the sound speed and density at the reference frequency, are found in a least-square sense at each spatial location within the medium.

Random variations in tissue parameters within the breast model, described in Sec. IV A, emulate real tissue variations that are unresolvable in magnetic resonance images. These random variations require different models of compressibility at each sample in the tissue model. Producing these models would be cost-prohibitive without an efficient method for mapping material properties to coefficients κ_ν and κ_∞ at every sample point. Although the previously noted scaling method is efficient, agreement between the computed attenuation slope, density, and sound speed with the random values assigned to each tissue type is not guaranteed at every spatial location.

An efficient but more accurate method was developed to ensure that coefficients κ_ν and κ_∞ can be computed at each spatial location to yield values of sound speed and attenuation that agree with reported measurement values at a desired reference frequency. The method uses a rational-function fit to optimal coefficients for a reference set of density, attenuation, and sound-speed values selected from the range of values represented in the medium. The following steps were used to determine the rational-function fit:

- (1) Optimal relaxation times τ_ν were calculated using a reference medium similar to body-temperature water, with a density of 1000 kg/m³, an acoustic absorption of 0.5 dB/cm/MHz, and a sound speed of 1524 m/s.
- (2) Unconstrained optimization was used to fit the coefficients κ_ν and κ_∞ calculated on a three-dimensional grid of density, attenuation, and sound speed values that span the range of values represented in the medium.
- (3) A variety of function forms were used to find, by trial and error, a low-order rational expression to yield an accurate representation of the coefficients κ_ν and κ_∞ . The selected rational function form was

$$\kappa_\nu(\beta, c_L, \rho) = \frac{1}{\rho} \left[\sum_{i=0}^4 \left(\frac{a_{\nu,i}}{c_L} + b_{\nu,i} \right) \beta^{i/2-1} \right], \quad (18)$$

where the coefficients $a_{\nu,i}$ and $b_{\nu,i}$ are empirically determined. The quantity c_L in Eq. (18) represents the sound

speed in the long-wavelength limit and is related to the coefficient κ_∞ by the expression²⁹

$$\kappa_\infty(\mathbf{r}) = \frac{1}{c_L^2(\mathbf{r})\rho(\mathbf{r})} - \sum_{\nu} \kappa_{\nu}. \quad (19)$$

Therefore, c_L must be determined everywhere so that the sound speed associated with the compressibility (17) agrees with the reference sound speed distribution at the desired frequency.

The rational expression (18) is an improvement over previous expressions that simply scale the compressibility amplitudes. The resulting efficient mapping facilitates the solution of an optimization problem at each spatial location to identify coefficients κ_ν and κ_∞ that yield desired values of sound speed and attenuation used in single-frequency solutions.

C. Techniques for massively parallel computing

The solution of large-scale wave propagation and scattering problems requires greater resources than are offered by a single desktop computer. Modern large-scale computing systems are typically composed of many networked, small nodes that each contain multiple processors in a shared-memory configuration. Although each processor may be treated as an independent system, higher computational efficiency may be realized if the hybrid shared and distributed nature of the computer system is leveraged to reduce inter-process communication. A distributed-memory implementation of the FMM³⁰ has been extended by using message passing to synchronize a group of multithreaded tasks running on multiple, distributed-memory nodes. Each task, however, may execute its threads on more than one CPU without the need to exchange messages. This configuration facilitates use of more parallel processors, with greater computational efficiency, than either a pure multithreaded or a pure message-passing implementation.

The distributed-memory FMM that forms the foundation of the hybrid implementation was originally designed to simulate electromagnetic scattering from perfectly conducting surfaces. For such problems, the computational complexity of the FMM is known to scale like $O(N \log N)$ rather than $O(N)$ because the recursive grouping of scattering elements balances a reduction in the number of groups at coarser levels in the hierarchy with a corresponding increase in the spatial-harmonic bandwidth of their associated plane-wave expansions.¹⁶ This means that communication and computational demands in a distributed-memory implementation remain significant at all levels of the FMM hierarchy. The authors of Ref. 30 developed the concept of *shared levels* to efficiently employ large numbers of parallel processes in the coarsest FMM levels. In a shared FMM level, rather than assign each interaction group wholly to a single process, each process is responsible for a subset of the plane-wave expansions for all groups in the level.

For the calculation of scattering (either acoustic or electromagnetic) by penetrable volumes, shared levels introduce unnecessary complexity. In a volume scatterer, the number

of groups in each successively coarser level of the FMM hierarchy diminishes faster than the increase in the harmonic bandwidth of the plane-wave expansions of the groups. Therefore, the amount of computation in coarser levels is less than that in the finer levels. As a result, the overall parallel efficiency of a distributed-memory FMM is much less sensitive to workload distribution and communication costs in the coarser levels. The sensitivity is further reduced in a hybrid shared- and distributed-memory implementation, where greater numbers of processors can be employed with no increase in communication costs.

Parallel efficiency, and the related concept of parallel speed-up, are important characteristics of parallel algorithms. On parallel systems maintained by large institutions for many users, high parallel efficiency not only corresponds to reduced calculation times, but is also important to justify the allocation of a portion of finite and in-demand resources to the solution of a particular problem. Parallel speed-up P is defined as $P = T_R/T_N$, where T_N is the time required to obtain a solution using N tasks (either distributed-memory processors or shared-memory threads) and T_R is the time required to obtain a solution using the minimum required number of tasks R . The parallel efficiency E is given by $E = RP/N$, which provides a measure of the fraction of computation time lost to synchronization or communication among tasks. In some applications, R must exceed the ideal value of unity because the computational or storage requirements of the calculation exceed the capabilities of a single task.

In the hybrid FMM, the principal cause for parallel-efficiency loss among many distributed-memory processes is, naturally, the need for communication among the processes. Asynchronous communication capabilities on massively parallel systems, which allow transmission and reception of messages to occur while each process is performing local work, can help mask communication latency and improve efficiency. The principal cause for efficiency loss among many threads in a single process is the existence of serialization points. For example, communication among multiple processes is always confined to a single thread, regardless of the number of threads available to each task. In addition, aggregation or distribution of plane-wave patterns is constrained to occur in ascending or descending level order, respectively. A thread that finishes its workload on one level of the FMM hierarchy cannot proceed to another level ahead of its peers.

The implementation of the k -space method, in which the dominant computational cost is that of a three-dimensional fast Fourier transform, employs a message-passing, distributed-memory version of the FFTW library that would require extensive modification to support hybrid shared- and distributed-memory paradigms. As a consequence, the implementation of the k -space method is a pure message-passing application. Because the parallel Fourier transform algorithm requires extensive communication between computing nodes, parallel acceleration is sensitive to node interconnect latency and throughput. The method therefore benefits from high-speed node interconnects employed on modern supercomputers.

IV. NUMERICAL RESULTS

Two numerical scattering experiments were conducted using three-dimensional tissue models derived from segmented magnetic resonance images. In the first experiment, a variety of machine configurations were employed in calculating scattering from a large-scale breast model to study the parallel efficiency and scalability of the FMM. These studies highlight programmatic compromises and suggest suitable parallel configurations for realistic computations. In the second experiment, single-frequency scattering of a 3.0-MHz, point-source field by a breast model of even larger scale was computed using the FMM. Comparable time-domain solutions were obtained from the k -space method and Fourier transformed to yield single-frequency pressure measurements that correspond to the FMM simulations.

A. Segmentation of magnetic resonance images

Studies of ultrasonic breast imaging require development of digitized breast models that can provide suitable foundations for realistic simulations of wave propagation, scattering and pulse-echo measurements. To model transmissions with a center frequency of 5 MHz and a 50% bandwidth, the tissue must be represented by three-dimensional arrays of mechanical parameters (sound speed, density and attenuation) that are sampled at intervals of about 50 μm . Thus, approximately 8×10^6 samples of each parameter are required to represent each cubic centimeter of tissue. Models of this scale were developed for this study from high-resolution MRI scans that were acquired from samples of donated post-mortem breast tissue. To image these samples, an excised volume of tissue surrounding the nipple was mounted in a special cylindrical container. The tissue in this container was surrounded by tubes filled with water and oil to facilitate separation of water and fat responses. Because scanning large volumes of breast tissue at a resolution of 50 μm is prohibitively costly and time consuming, the models were developed from scans acquired on a grid of approximately $350 \times 300 \times 150$ voxels with an isotropic resolution of 200 μm . Although 200- μm resolution is higher than that of typical MRI scans, the desired model has a resolution four times finer. The 200- μm resolution is sufficient, however, because the MRI data are only used to segment primary tissue types. Fine structure that simulates the effect of tissue histology on wave propagation was generated randomly by parametric procedures using different parameters for each tissue type.

The MRI data were segmented into distinct tissue types using a procedure that differs from a prior attempt⁴ that only characterized connective tissue. Four spoiled-gradient (SPGR) scans were acquired with different echo times along with a T2-weighted scan. These protocols are standard MRI scanning procedures that are used to illicit responses to different tissue characteristics.³¹ The SPGR scans were used to separate water and fat responses, while the T2 scans were used to highlight ductal structure. There were also differences in the frame-to-frame separations between the SPGR and T2 scans caused by small variations in the field gradients used to isolate local responses. Despite these variations, a

FIG. 1. (Color online) T2 response levels (left) are used to separate the water response into three tissue types (large and extensive ductal structures, small and sparse ductal structures, and dense connective tissue) that are merged with the fat segmentation (right). The white line has length 1 cm.

mapping between the SPGR and T2 frames yielded close agreement between the outer boundaries and internal structures of the SPGR and T2 images.

Separation of water and fat responses in SPGR images was employed to distinguish in the fat response clusters of adipose cells with diameters smaller than a millimeter. Connective tissues encapsulating larger fat lobules (including Cooper's ligaments, internal septae and the sheets of connective tissue that separate the mammary zone from the pre-mammary and retro-mammary zones) were interpreted as the boundaries of regions with strong fat responses. This method of segmenting fascia gives better results than using the boundaries of regions of strong water responses because partial volume effects from the thin sheets of connective tissue obscure the gradients in water responses more than they obscure the gradients in fat responses.

Remaining tissues in the regions where the water response is dominant are identified by distinguishing three levels of the T2 response. Because strong T2 responses generally indicate the presence of free fluid, the strongest T2 responses are characterized as large ducts or blood vessels. However, some of these regions are too spread out to be individual ducts and are interpreted as ductal networks that are too fine to be resolved. The fluid in regions with intermediate T2 response levels also implies that the regions contain ductal structure. However, reduced levels of the T2 response suggest smaller individual ducts or less extensive ductal networks. Low T2 response levels are interpreted as the foundation of dense connective tissue in which the glandular tissues are embedded. As each region is extracted, morphological smoothing operations are applied to produce binary tissue masks that are not pixelated. A thin layer of epidermal tissue, which is not distinguished in the MRI scans but is evident in microscope images of corresponding slices of the tissue sample, is assigned to the voxels along the top surface of the sample. The image in Fig. 1 shows the T2 segmentation combined with the SPGR fat segmentation.

The scattering model was extended beyond the breast specimen by grafting the high-resolution segmentation onto a lower-resolution model of the chest wall derived from an *in vivo* MRI scan of a different subject. The low-resolution scan was T1 weighted without fat suppression. A frame of the grafted model is shown in Fig. 2. Fat, connective tissue, and pectoral muscle were segmented from the low-resolution scan and then interpolated to the 200- μ m resolution of the *ex vivo* scan.

A complex transformation was used to reshape the segmented, high-resolution model to mate smoothly with the segmented tissues from the *in vivo* scan. The transformation used to reshape the segmented model was developed from the complex function $w = \arccos z$. This function implicitly generates a two-dimensional transformation $(x, y) \mapsto (X, Y)$ by the assignments $z = xe^{iy}$ and $w = Xe^{iY}$. The transformation was applied to the breast model by assigning a coordinate pair (x, y) to each voxel that specifies, respectively, the height of the voxel above the chest wall and the lateral offset of the voxel from the nipple. The transformed pair (X, Y) gives revised coordinates that reposition each voxel. The deformation induced by the transformation partially restores the natural shape of the tissue prior to surgical excision. To avoid unnatural truncation of internal structures along the boundary between the two models, the connective tissue of each model was allowed to cross the boundary into the other model. In addition, surface tissues of the high-resolution segmentation were smoothly extended, where necessary, to join with the skin at the outer edges of the low resolution model.

The resulting segmented model, interpolated to 50- μ m resolution on a grid measuring $1400 \times 1200 \times 840$ voxels, was reviewed by a breast ultrasound clinician, a pathologist and a breast surgeon to ensure consistency with the anatomy and histology of the modeled tissue. The model was

FIG. 2. (Color online) Graft of the high-resolution breast model onto a low-resolution model from an *in vivo* scan of a different subject. The joint between the models is depicted with a dotted line. The white line has length cm.

truncated to $1350 \times 1170 \times 790$ voxels by eliminating planes along every face of the imaging domain that contained boundary artifacts from the magnetic resonance image or that contained no breast tissue.

Average values for the tissue mechanical properties of the seven segmented tissue types were adapted from published measurements.^{32,33} These values, as measured at 3.0 MHz, are listed in the first three columns of Table I. However, the segmentation procedure separates gland into three different tissues based on T2 response levels that are listed in the table as dense connective tissue; small, sparse duct; and large, dense duct. Published tissue properties for glandular tissues do not make these distinctions. Thus, the average mechanical properties for all three components of glandular tissue were chosen to be roughly the same. Small adjustments to the properties of large, dense duct were adopted to make these values more consistent with the ductal properties reported in Ref. 32.

Random variations in material properties were produced from a single distribution of point scatterers with a fractional density of 0.2 that extends across the entire volume of tissue. This resulted in scattering centers at approximately one fifth of the voxels in the volume. This distribution of point scatterers was convolved with a spatial Gaussian with a standard deviation of 0.03 mm to endow the point scatterers with a structure function of non-infinitesimal size. The value of a specific tissue property P (sound speed, attenuation, or density) in a particular tissue type took the form

$$P(\mathbf{r}) = M(\mathbf{r})\langle P \rangle \left(1 + \left[\frac{R}{100} \right] \left[\frac{S(\mathbf{r}) - \langle S \rangle}{\sqrt{\langle S^2 \rangle}} \right] \right), \quad (20)$$

where $\langle \cdot \rangle$ denotes averaging, $\langle P \rangle$ represents the average value for the tissue property listed in Table I, R is the percentage of root-mean-square variation for the tissue property listed in Table I, S denotes the distribution of random sources described above, and M denotes a smoothed binary mask for the tissue that is produced by convolving the tissue mask with a Gaussian filter with a standard deviation of 0.25 mm. Smoothing of the binary mask softens abrupt boundaries between tissues that could cause unnatural specular reflection in simulations of ultrasonic propagation.

The size (0.03 mm) and fractional density (0.2) of the scatterers were selected based on internal, unpublished studies of the speckle patterns in b-scan images of simulated

tissue samples, and also internal, unpublished studies of forward propagation through the same tissue samples. The speckle intensity was found to be strongest for scatterers that are roughly one eighth of a wavelength at the b-scan center frequency. At a frequency of 5 MHz, one eighth of a wavelength is near 0.03 mm. Actual breast tissue has random scatterers of many different shapes and sizes, but statistics of these scattering distributions are not available. For this reason, random variations in the model breast were based on the simplified model, described above, of single-sized Gaussian scatterers with size and density parameters that result in realistic b-scan responses.

Amplitudes of the random variations were also determined from b-scan simulations. The echogenicity of the simulated tissues was adjusted by changing the percentage of random variation listed in Table I. A nominal value of 2% root-mean-square variations (the value used in Ref. 32) was increased or decreased to produce sonogram responses that were consistent with the expected echogenicities of different tissues. The b-scan simulations were derived exclusively from variations in sound speed, but the amounts of random variation developed from these simulations were also used in the models of attenuation and density.

The effect of the tissue model statistics on forward propagation was assessed by computing the phase aberration of plane-wave transmissions through slabs of random tissue composed of scatterers with sizes of 0.01, 0.03, and 0.05 mm, and with random sound speed variations with RMS values ranging from 0.5% to 3.0% of the average sound speed. The aberration produced by 0.03-mm scatterers with RMS variations equal to 2% of the average sound speed were found to be consistent with the expected aberration from breast tissue, confirming the selection of these parameters based on b-scan simulations.

B. Scaling studies

The Kraken supercomputer system was employed to study the scalability of the FMM solver. Maintained by the National Institute for Computational Sciences at the University of Tennessee, Kraken is a Cray XT5 composed of 9408 compute nodes, each containing 2, 6-core processors and 16 GB of RAM. The nodes of Kraken are connected by a Cray SeaStar2+ high-speed network. Comparable studies of parallel scaling of the k -space method were previously performed^{20,21} on the Cray XT-4 system Franklin, maintained by the United States Department of Energy as part of the National Energy Research Scientific Computing Center (NERSC). These prior scaling studies investigated problems with scales comparable to those described here and confirmed a parallel efficiency of approximately 90% for problem sizes six times larger than the problem of interest when using 2048 CPU cores interconnected by a less-capable network.

The previously described $1350 \times 1170 \times 790$ -voxel breast model was unsuitable for scaling studies because the large problem scale prohibited repeated solutions. Instead, magnetic resonance images of a separate $59 \text{ mm} \times 59.5 \text{ mm} \times 14.5 \text{ mm}$ breast specimen were segmented and converted

TABLE I. Average properties (measured at 3 MHz) and variation in segmented tissues.

Tissue	Speed (mm/ μ s)	Density (rel. H ₂ O)	Attenuation (dB/cm/MHz)	Variation (pct.)
Fat	1.47	0.95	0.36	1.5
Fascia	1.55	1.12	1.15	2.5
Connective	1.55	1.06	0.72	2.5
Sparse duct	1.55	1.06	0.72	1.5
Dense duct	1.60	1.04	0.91	0.5
Skin	1.59	1.09	1.15	2.5
Muscle	1.54	1.05	2.00	0.8

into a $1180 \times 1190 \times 290$ -voxel numerical model using the aforementioned procedure. The resulting model was small enough to facilitate repeated solutions, but still complex and large enough to yield meaningful scaling results. Calculations of scattering of a 3.0-MHz, point-source field by this model, which requires more than 400×10^6 unknowns, were performed using the hybrid FMM. Run times were measured for three types of configurations:

- (1) A variable number of message-passing tasks, each using a fixed number of threads;
- (2) A fixed number of message-passing tasks, each using a variable number of threads;
- (3) A variable number of both message-passing tasks and threads such that the product of threads and tasks (the total number of CPUs employed) remained constant.

Although cubic domains are simple to generate, the symmetry of the grid can mask inefficiencies that might otherwise appear. Because the breast has a relatively narrow edge in the third dimension, the FMM hierarchy is not optimally filled, which can present challenges during the aggregation and distribution stages of computation. These challenges help ensure that observed scaling results reflect real-world performance of the methods.

The large computational size of the scattering problem required the memory resources of at least 8 Kraken nodes. Therefore, rather than comparing performance of multiple message-passing tasks to the performance of an ideal single task, the parallel speed-up and efficiency were measured relative to the minimum number of tasks required of the simulation.

For a fixed number of shared-memory threads per task, the parallel speed-up and efficiency are shown as functions of the number of message-passing tasks in the upper panel of Fig. 3. The “TL2” labels refer to configurations in which the coarsest FMM computations occur in the second level of the FMM hierarchy, where the domain is subdivided into 4 groups per dimension, or a total of 64 groups; “TL3” labels refer to configurations in which the coarsest FMM computations occur in the third level of the FMM hierarchy, where the domain is subdivided into 8 groups per dimension, or a total of 512 groups. The TL3 study is significant because it provides insight into potential losses in efficiency when the number of groups at the coarsest level is less than the number of MPI tasks assigned to computation. In both configurations, the seventh level grouping (a total of 2097152 groups, or 128 per dimension) was the finest subdivision used in the FMM. Each parallel task employed six shared-memory threads. On Kraken, the minimum number of tasks was $R=8$, which corresponds to one task per node on each required node. The “process count multiplier” label refers to the ratio of the chosen number of tasks to R .

Parallel speed-up and efficiency are shown as functions of the number of threads per task in the lower panel of Fig. 3. For these experiments, 64 message-passing tasks were employed. The reference number of threads was $R=1$.

As expected, the parallel efficiency declines as either the number of message-passing tasks or the number of threads per task is increased. Hence, the optimum

FIG. 3. (Color online) Parallel speed-up and efficiency for a large-scale FMM scattering solution as functions of (top) the number of message-passing tasks, each using six shared-memory threads; and (bottom) the number of shared-memory threads employed by each of 64 message-passing tasks.

configuration in terms of parallel efficiency would be to employ the minimum possible number of tasks, each running a single computation thread. However, while this may be most efficient in terms of resource usage, it does not offer the fastest solution. An alternative measure of optimality is obtained by fixing the total number of processors enlisted to solve a scattering problem, and varying the proportion of threads per task and total tasks to identify the minimum wall time for a solution. These observations are recorded in Fig. 4, where a total of 240 processors were employed to compute 15 GMRES iterations of a scattering solution. The numbers of tasks and threads per task were varied so the product of the two remained constant. In this example, the optimum time was observed when 40 message-passing tasks were employed, each running 6 threads. This arrangement coincides with the hardware design, where each 6-core processor was devoted to a single message-passing task.

FIG. 4. (Color online) Wall time for 15 GMRES iterations of a large-scale FMM acoustic solver as a function of the number of threads per task and the total number of message-passing tasks. The product of the numbers of tasks and threads was constant.

C. Comparison of solution methods

Comparisons of the FMM solution to a corresponding k -space solution were performed by placing the previously described $1350 \times 1170 \times 790$ -voxel breast model such that its center coincided with the coordinate origin and the nipple extended in the $+z$ direction. The tissue model included variations in sound speed and attenuation but not in density. A point source was positioned at $(x, y, z) = (0, 0, 30)$ mm, approximately 10 mm from the $+z$ face of the computational domain. In the FMM simulation, the source was excited at 3.0 MHz. At 3.0 MHz, the breast model is sampled at approximately ten points per wavelength and the FMM implementation converged rapidly to a solution; increasing the frequency to 5.0 MHz (to agree with time-domain solutions) reduced this sampling to near six points per wavelength and a convergent solution could not reasonably be obtained. Versions of the model with and without acoustic absorption were considered.

In the k -space simulation, the sources were excited with a wideband pulse with a center frequency of 5.0 MHz and a -6 -dB bandwidth of 2.5 MHz. These parameters coincide with those that are typically employed in clinical ultrasound systems used for diagnostic imaging. The pulse was sampled by 9550 time steps, each separated by $\Delta t = 8.33$ ns (which corresponds to a Courant-Friedrichs-Lewy number of approximately 0.3). This number of steps allowed the pulse to travel from the point in the computational domain closest to the source location to the farthest point in the domain and back. The time-domain amplitude of the scattered pressure in the center of the computational domain was found to be more than 100 dB below its peak value after executing the time steps. The breast tissue model had a spatial resolution sufficient for Nyquist sampling up to 7.0 MHz, which is effectively the highest frequency in the pulse: The spectral amplitude of the incident pulse at 7.0 MHz was more than 80 dB below the amplitude at the 5.0-MHz center frequency. In accordance with Ref. 25, the medium was smoothed using half-band filtering to limit the spatial-frequency bandwidth of the medium variations and improve the accuracy of backward scattering. A 10-cell PML with an attenuation profile described in Ref. 24 was used to truncate the computational domain in the k -space simulation.

The computational requirements of a problem of such large scale, consisting of substantial computer memory and intensive calculations to produce both the FMM and k -space results involving more than 1.2×10^9 unknowns, were satisfied by performing the calculations on the previously described Kraken supercomputer. The k -space simulation, which required 54 h, was executed on 816 CPU cores distributed among 68 compute nodes. The 3.0-MHz FMM simulation required approximately 1.5 h using 3072 CPU cores distributed among 256 compute nodes.

The large number of temporal samples in the k -space simulations prohibited comparisons of the field at all points in the computational domain. As a consequence, the time-domain pressures within several representative planes were stored and subsequently Fourier transformed to obtain the single-frequency pressure fields at 3.0 MHz for comparison

to the FMM results. Comparisons were made in the $x=0$, $y=0$, and $z=0$ planes for both the lossless and lossy models. In addition, comparisons of the lossy model were made in 97 constant- z planes evenly distributed throughout the model volume.

The single-frequency FMM solution was taken as a reference for comparisons because the FMM-calculated fields were available at every point within the model volume. The normalized, per-plane (or per-slice), root-mean-square error

$$E_P = \sqrt{\frac{\int_P d\mathbf{r} |p_f(\mathbf{r}) - p_k(\mathbf{r})|^2}{\int_D d\mathbf{r} |p_f(\mathbf{r})|^2 / N_P}} \quad (21)$$

was computed for fields in each plane, where P is the plane in which the pressure has been sampled, D is the entire model volume, N_P is the total number of model samples along the axis perpendicular to the plane, p_f is the total pressure computed using the FMM, and p_k is the total pressure computed using the k -space method. The quantity in Eq. (21) corresponds, in a sense, to the fraction of the expected energy in each plane corresponding to differences between the two solutions. A corresponding volume root-mean-square error can be obtained by integrating the error E_P in the RMS sense,

$$E_V = \sqrt{\frac{\int_D d\mathbf{r} |p_f(\mathbf{r}) - p_k(\mathbf{r})|^2}{\int_D d\mathbf{r} |p_f(\mathbf{r})|^2}} = \sqrt{\frac{1}{N_P} \int dP E_P^2}, \quad (22)$$

where the second integral in Eq. (22) is interpreted as a summation over a set of parallel sample planes that span the model volume D . The per-slice RMS errors are listed in Table II for representative planes. The volume RMS error for the lossy comparison, calculated using Eq. (22) for the 97 constant- z planes distributed throughout the volume, was 2.6%.

For sources placed a finite distance away from the model volume, Eq. (21) tends to underestimate error further from the source because the average per-slice energy does not account for propagation of energy beyond the bounds of the model. An alternative per-plane RMS error, given by the expression

$$E'_P = \sqrt{\frac{\int_P d\mathbf{r} |p_f(\mathbf{r}) - p_k(\mathbf{r})|^2}{\int_P d\mathbf{r} |p_i(\mathbf{r})|^2}}, \quad (23)$$

TABLE II. Per-slice errors between k -space and FMM scattering.

Plane	RMS Error (%)	
	Lossless	Lossy
$x=0$	3.0	3.5
$y=0$	3.1	3.8
$z=0$	2.4	2.8

FIG. 5. (Color online) Error in the k -space solution of scattering by a lossy breast specimen relative to a 3-MHz FMM solution. Slices are perpendicular to the z axis.

uses as a reference the energy of the incident field p_i in the plane P to account for this loss of energy. While the adjusted measure (23) emphasizes error more uniformly throughout the model volume, the measure may not readily be used to define a volume error that properly accounts for absorptive losses. A comparison of the two per-slice error measures is shown in Fig. 5 for the lossy model.

In analogy to the per-slice error, a normalized error contribution C may be defined by

$$C(\mathbf{r}) = \frac{|p_f(\mathbf{r}) - p_k(\mathbf{r})|}{\sqrt{\int_D d\mathbf{r} |p_f(\mathbf{r})|^2 / (N_C N_R)}}, \quad (24)$$

where N_C and N_R represent the number of columns and rows, respectively, in a slice P of the model volume D . The error contributions for representative planes of the lossy model are shown in Fig. 6. The scattered field amplitude in each of these planes, relative to the peak amplitude of the incident field within the computational domain, is shown in Fig. 7 as computed by the FMM. The error contribution for the

FIG. 6. Error contribution in the $x=0$, $y=0$, and $z=0$ planes for FMM and k -space calculations of scattering of a 3.0-MHz point-source field by the lossy segmented breast model.

FIG. 7. Scattered field amplitude, relative to the peak amplitude of the incident field in the computational domain, in the $x=0$, $y=0$, and $z=0$ planes for the FMM calculation of scattering of a 3.0-MHz point-source field by the lossy segmented breast model.

lossless case has a similar structure with lower peak values. The results in the figure suggest that no obvious relation exists between position and error, other than a general concentration of error within the breast tissue. Specifically, the quality of the time-domain solution does not degrade as the wave propagates through the breast model in the z direction. The increased error observed in Fig. 5 along the propagation direction is instead attributable to the conical shape of the breast; there is simply more scattering tissue (hence, more points at which the fields computed by the two methods disagree) to contribute to the per-slice error further along the propagation axis.

V. CONCLUSION

Segmentation of magnetic resonance images of *ex vivo* breast specimens provides a classification of tissue types that has been used to develop numerical models of sound speed, attenuation and density in the specimen. The models include random fluctuations to account for unresolved scattering features. The large scale of the resulting model presents computational challenges in the simulation of ultrasound propagation and scattering that are addressed by massively parallel implementations of the fast multipole and k -space methods. A hybrid shared- and distributed-memory parallelization scheme was presented for the FMM that can mitigate rapid decreases in parallel efficiency as the number of CPUs increases. The effectiveness of the hybrid design is demonstrated by a scaling study involving a fixed number of CPUs in configurations that varied the number of threads and distributed-memory tasks divided among the CPUs. Selecting an optimum configuration for a fixed number of CPUs improves overall parallel efficiency that was observed to reduce wall time by up to 25%. The k -space method, which relies on a well-studied and thoroughly optimized

distributed-memory FFT library, was previously found to exhibit excellent parallel scaling for large numbers of CPUs without the need for threading. Both the FMM and k -space method, each previously validated against analytic results for simple spherical structures, were employed to compute scattering by the numerical breast model, with and without absorption, of the field generated by a point source. Comparisons of the solutions at 3.0 MHz (the frequency at which the sampling density of the medium is roughly ten points per wavelength per dimension) in representative planes revealed differences less than 3% RMS. The close agreement of these disparate methods establishes confidence in the accuracy of each for the solution of ultrasound scattering problems for which true solutions are unavailable.

ACKNOWLEDGMENTS

This research was funded in part by grants R01 EB 009692 and R01 EB 010069 from the National Institute of Biomedical Imaging and Bioengineering of the National Institutes of Health. Calculations associated with this work used resources of the National Energy Research Scientific Computing Center, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231; and the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number OCI-1053575. The authors thank Professor Weng C. Chew at the University of Illinois for the use of his ScaleME FMM library (see Refs. 16 and 30). The authors also thank Dr. David G. Hicks, Dr. Howard N. Langstein, and Dr. Avicé M. O'Connell, all at the University of Rochester Medical Center, for their extensive consultation while the segmented model was refined to conform to expected morphology. Finally, the authors thank Dr. Joseph A. Sperryak at the Roswell Park Cancer Institute for designing and executing magnetic resonance imaging procedures used to produce the breast model.

- ¹M. Tabei, T. D. Mast, and R. C. Waag, "Simulation of ultrasonic focus aberration and correction through human tissue," *J. Acoust. Soc. Am.* **113**, 1166–1176 (2003).
- ²J. C. Lacefield, W. C. Pilkington, and R. C. Waag, "Distributed aberrators for emulation of pulse distortion by abdominal wall," *Acoust. Res. Lett. Online* **3**, 47–52 (2002).
- ³T. D. Mast, L. M. Hinkelman, L. A. Metlay, M. J. Orr, and R. C. Waag, "Simulation of ultrasonic pulse propagation, distortion, and attenuation in the human chest wall," *J. Acoust. Soc. Am.* **106**, 3665–3677 (1999).
- ⁴G. S. Salahura, J. C. Tillett, L. A. Metlay, and R. C. Waag, "Large-scale propagation of ultrasound in a 3-D breast model based on high-resolution MRI data," *IEEE Trans. Biomed. Eng.* **57**, 1273–1284 (2010).
- ⁵R. C. Waag and J. P. Astheimer, "Statistical estimation of ultrasonic propagation path parameters for aberration correction," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control* **52**, 851–869 (2005).
- ⁶J. C. Tillett, J. P. Astheimer, and R. C. Waag, "A model of distributed phase aberration for deblurring phase estimated from scattering," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control* **57**, 214–228 (2010).
- ⁷A. J. Devaney and M. L. Oristaglio, "Inversion procedure for inverse scattering within the distorted-wave Born approximation," *Phys. Rev. Lett.* **51**, 237–240 (1983).

- ⁸W. C. Chew and Y. M. Wang, "Reconstruction of two-dimensional permittivity distribution using the distorted Born iterative method," *IEEE Trans. Med. Imag.* **9**, 218–225 (1990).
- ⁹R. Lavarello and M. Oelze, "Density imaging using inverse scattering," *J. Acoust. Soc. Am.* **125**, 793–802 (2009).
- ¹⁰A. J. Hesford and W. C. Chew, "Fast inverse scattering solutions using the distorted Born iterative method and the multilevel fast multipole algorithm," *J. Acoust. Soc. Am.* **128**, 679–690 (2010).
- ¹¹T. D. Mast, A. I. Nachman, and R. C. Waag, "Focusing and imaging using eigenfunctions of the scattering operator," *J. Acoust. Soc. Am.* **102**, 715–726 (1997).
- ¹²F. Lin, A. I. Nachman, and R. C. Waag, "Quantitative imaging using a time-domain eigenfunction method," *J. Acoust. Soc. Am.* **108**, 899–912 (2000).
- ¹³R. C. Waag, F. Lin, T. K. Varslot, and J. P. Astheimer, "An eigenfunction method for reconstruction of large-scale and high-contrast objects," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control* **54**, 1316–1332 (2007).
- ¹⁴L. Greengard and V. Rokhlin, "A fast algorithm for particle simulations," *J. Comput. Phys.* **73**, 325–348 (1987).
- ¹⁵V. Rokhlin, "Rapid solution of integral equations of scattering theory in two dimensions," *J. Comput. Phys.* **86**, 414–439 (1990).
- ¹⁶*Fast and Efficient Algorithms in Computational Electromagnetics*, edited by W. C. Chew, J. Jin, E. Michielssen, and J. Song (Artech House, Boston, MA, 2001), pp. 39–150.
- ¹⁷N. N. Bojarski, "The k -space formulation of the scattering problem in the time domain," *J. Acoust. Soc. Am.* **72**, 570–584 (1982).
- ¹⁸N. 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).
- ¹⁹B. Compani-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).
- ²⁰D. A. Donzis, P. K. Yeung, and D. Pekurovsky, "Turbulence simulations on $O(10^4)$ processors," in *Proc. TeraGrid'08* (Las Vegas, NV, 2008).
- ²¹J. C. Tillett, M. I. Daoud, J. C. Lacefield, and R. C. Waag, "A k -space method for acoustic propagation using coupled first-order equations in three dimensions," *J. Acoust. Soc. Am.* **126**, 1231–1244 (2009).
- ²²A. J. Hesford and R. C. Waag, "The fast multipole method and Fourier convolution for the solution of acoustic scattering on regular volumetric grids," *J. Comput. Phys.* **229**, 8199–8210 (2010).
- ²³A. J. Hesford and R. C. Waag, "Reduced-rank approximations to the far-field transform in the gridded fast multipole method," *J. Comput. Phys.* **230**, 3656–3667 (2011).
- ²⁴M. 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).
- ²⁵T. 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).
- ²⁶Y. Saad and M. H. Schultz, "GMRES—a generalized minimal residual algorithm for solving nonsymmetric linear-systems," *SIAM J. Sci. Stat. Comput.* **7**, 856–869 (1986).
- ²⁷H. A. van der Vorst, "Bi-CGSTAB: A fast and smoothly converging variant of Bi-CG for the solution of nonsymmetric linear systems," *SIAM J. Sci. Stat. Comput.* **13**, 631–644 (1992).
- ²⁸J. P. Berenger, "A perfectly matched layer for the absorption of electromagnetic waves," *J. Comput. Phys.* **114**, 185–200 (1994).
- ²⁹A. I. Nachman, J. F. Smith III, and R. C. Waag, "An equation for acoustic propagation in inhomogeneous media with relaxation losses," *J. Acoust. Soc. Am.* **88**, 1584–1594 (1990).
- ³⁰S. Velampambal and W. C. Chew, "Analysis and performance of a distributed memory multilevel fast multipole algorithm," *IEEE Trans. Antennas Propag.* **53**, 2719–2727 (2005).
- ³¹R. E. Hendrick, *Breast MRI: Fundamentals and Technical Aspects* (Springer, New York, 2008), pp. 135–170.
- ³²E. Franceschini, S. Mensah, D. Amy, and J.-P. Lefebvre, "A 2-D anatomic breast ductal computer phantom for ultrasonic imaging," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control* **53**, 1281–1288 (2006).
- ³³E. L. Madsen, J. A. Zagzebski, G. R. Frank, J. F. Greenleaf, and P. L. Carson, "Anthropomorphic breast phantoms for assessing ultrasonic imaging system performance and for training ultrasonographers: part I," *J. Clin. Ultrasound* **10**, 67–75 (1982).