Journal of Computational Physics 229 (2010) 8199–8210

Contents lists available at ScienceDirect

Journal of Computational Physics

journal homepage: www.elsevier.com/locate/jcp

The fast multipole method and Fourier convolution for the solution of acoustic scattering on regular volumetric grids

Andrew J. Hesford [a,*], Robert C. Waag [a,b]

[a] Department of Electrical and Computer Engineering, University of Rochester, Rochester, NY 14642-8648, USA
[b] Department of Imaging Sciences, University of Rochester, Rochester, NY 14642-8648, USA

ARTICLE INFO

Article history:
Received 25 February 2010
Received in revised form 19 May 2010
Accepted 21 July 2010
Available online 4 August 2010

Keywords:
Fast solvers
Iterative methods
Acoustic scattering
Fast multipole method
Moment methods
Fast Fourier transform

ABSTRACT

The fast multipole method (FMM) is applied to the solution of large-scale, three-dimensional acoustic scattering problems involving inhomogeneous objects defined on a regular grid. The grid arrangement is especially well suited to applications in which the scattering geometry is not known *a priori* and is reconstructed on a regular grid using iterative inverse scattering algorithms or other imaging techniques. The regular structure of unknown scattering elements facilitates a dramatic reduction in the amount of storage and computation required for the FMM, both of which scale linearly with the number of scattering elements. In particular, the use of fast Fourier transforms to compute Green's function convolutions required for neighboring interactions lowers the often-significant cost of finest-level FMM computations and helps mitigate the dependence of FMM cost on finest-level box size. Numerical results demonstrate the efficiency of the composite method as the number of scattering elements in each finest-level box is increased.

1. Introduction

Acoustic scattering from large-scale, inhomogeneous, penetrable structures in three dimensions has a wide range of applications. For example, estimation and correction of aberration requires simulation of propagation of ultrasound pulses through aberrating media [1–3]. Iterative inverse scattering algorithms such as the distorted Born iterative method [4–7] and the eigenfunction method [8–10] require repeated computation of the forward scattering problem for successive guesses of a reconstructed medium. The arbitrary nature of the scattering media, the desire for automatic satisfaction of the radiation condition, and the embarrassingly parallel nature of solutions at multiple frequencies make frequency-domain integral equation solvers particularly attractive where repeated calculations are required.

The fast multipole method (FMM) [11–13] has emerged as a widely used algorithm [14] for efficient solution of integral equations that describe electromagnetic or acoustic scattering in two or three dimensions. Through recursive subdivision of the media into groups of scattering elements, the FMM computes the product of a test vector and the Green's function matrix in a time that is $O(N)$ when the medium is composed of N scattering elements distributed in a volume-filling fashion. Like methods such as the adaptive cross approximation [15–17], efficient solutions are obtained by eliminating approximately redundant information from interactions between sufficiently separated groups. However, these methods achieve a common goal using different techniques. The adaptive cross approximation uses an algebraic method to construct products of low-rank matrices that cast all pairwise interactions in terms of fewer, dominant interactions. In contrast, the FMM discards redundant information by representing interactions using band-limited multipole expansions.

* Corresponding author. Tel.: +1 585 275 6530.
E-mail address: hesford@ece.rochester.edu (A.J. Hesford).

0021-9991/$ - see front matter © 2010 Elsevier Inc. All rights reserved.
doi:10.1016/j.jcp.2010.07.025

In traditional implementations, the FMM has significant computational overhead and storage requirements (that both scale linearly with the problem size). Normally, the computational overhead is overlooked, because it can be amortized over the many matrix–vector products computed during iterative solution of the scattering problem. Another key issue associated with the FMM is its sensitivity to the size of finest-level scattering groups: when the groups are too small, the diagonal translation operators suffer from low-frequency breakdown and must be replaced with alternative operators. For very large groups, efficiency suffers because interactions between neighboring groups (that cannot be calculated using diagonal translation operators) are normally represented as dense Green's function matrices.

Regular, gridded arrangements of scattering elements in three dimensions arise in many problems of interest. In the aforementioned inverse methods, the scattering structure is not known *a priori*, rendering more sophisticated geometric meshes useless. Additionally, scattering models developed with other techniques, such as magnetic resonance or b-scan imaging, are often defined on a regular grid of voxels. Regular arrangements of unknowns offer simplicity of description and visualization.

When the scattering elements are arranged in a regular grid, variants of the conjugate gradient-fast Fourier transform (CG-FFT) method have also been applied to scattering problems in three dimensions [18–21]. These methods exploit the convolutional nature of the Green's function to evaluate the product of a test vector and the Green's function matrix as a Hadamard product in the spectral domain. Such methods have very little computational overhead and are easily understood and implemented. The comparatively large scaling constant and conceptual complexity of the FMM mean that CG-FFT methods may remain competitive for problems of intermediate scale.

The most substantial issues associated with the FMM can be avoided when scattering elements are defined on a regular grid. At the finest FMM level, neighboring interactions are directly computed as a convolution of the Green's function and a source pressure distribution. Elements defined on a grid allow this convolution to be accelerated using the FFT in the same fashion as CG-FFT methods. The reduced complexity afforded by FFT convolution diminishes the dependence of the FMM on the finest-level group size, allowing groups to grow large enough to avoid low-frequency breakdown or to match the natural structure of the scattering geometry (for example, a sparse arrangement of smaller scatterers occupying a large volume). Additionally, the regular arrangement of scattering elements, when made to coincide with the grid imposed by recursive FMM group subdivisions, can be exploited to reduce dramatically the computational and storage overhead associated with FMM setup. This reduction in setup overhead eliminates a key advantage held by CG-FFT methods over the FMM.

Two principal factors motivate the use of the FMM over pure CG-FFT methods even when scattering elements are defined on a regular grid. First, the asymptotic complexity of methods based on FFT convolution of the Green's function with a pressure distribution is $O(N\log N)$ for N scattering elements. The $O(N)$ asymptotic complexity of the FMM for volume scattering means that, in the large-problem limit, the FMM will eventually be more efficient than purely FFT-based methods. In three dimensions, even modest increases in problem scale can rapidly push N above this crossover point. Second, CG-FFT methods can not accommodate scattering geometries which contain large regions of homogeneous background material without modeling these non-scattering regions. With the FMM, gridded arrangements of scattering elements do not need to be connected, and the background material does not need to be modeled. This results in a reduction in the overall unknown count that can dramatically improve efficiency.

This paper presents an implementation of the fast multipole method that uses FFT convolution to represent neighboring interactions at the finest level and that exploits the regular arrangement of basis functions to reduce significantly the memory demands and setup overhead of the fast multipole method. The efficiency provided by regular arrangements of scattering elements in the FMM has been previously recognized [7,22], and the use of a hybrid FFT and FMM was presented in Ref. [22] for electromagnetic scattering from dielectric media. However, Ref. [22] was primarily interested in arrangements of sparse scatterers, with FFT convolution representing interactions within each entire scatterer. The FMM interactions in Ref. [22] were only used to convey interactions between distinct, whole scatterers. This is inefficient when the scatterers are large because of the asymptotic scaling of FFT convolution is worse than that of the FMM, and because of the cost of evaluating far-field FMM signatures becomes prohibitively expensive. The application of FFT convolution to neighboring interactions within a dense scatterer requires a novel formulation that was recognized in the noted work, but was not developed. This paper presents an explicit development of a hybrid FFT-FMM formulation that uses the FMM within dense scatterers and relies on FFT convolution only for near-field interactions. In addition, this paper explores the scaling of the FMM with FFT convolution as the numbers of elements in finest-level groups are varied significantly, thereby offering insight into issues related to large groups and offering techniques to avoid these issues.

2. Theory

Problems of interest involve three-dimensional scattering of acoustic waves in the presence of a bounded, penetrable, inhomogeneous scatterer with arbitrarily variable sound speed $c(\mathbf{r})$, attenuation $\alpha(\mathbf{r})$, and density $\rho(\mathbf{r})$, where $\mathbf{r}$ is a three-dimensional coordinate vector. The scatterer will be considered to be embedded in an infinite, homogeneous background medium with sound speed c_0, attenuation α_0, and density ρ_0. The complex wave number describing the scatterer is [10,23]

$$k(\mathbf{r}) = \frac{\omega}{c(\mathbf{r})} + i\alpha(\mathbf{r}), \tag{1}$$

where $\omega = 2\pi f$ is the radian frequency of the time-harmonic pressure fields. Assuming a time dependence of $e^{-i\omega t}$, the total acoustic pressure $p(\mathbf{r})$ satisfies the acoustic wave equation [24]

$$\rho(\mathbf{r})\nabla\cdot\left[\rho^{-1}(\mathbf{r})\nabla p(\mathbf{r})\right] + k^2(\mathbf{r})p(\mathbf{r}) = -S(\mathbf{r}), \tag{2}$$

where S is an arbitrary acoustic source distribution. With the change of variable $p(\mathbf{r}) = \rho_n^{1/2}(\mathbf{r})f(\mathbf{r})$, where $\rho_n = \rho/\rho_0$, the wave Eq. (2) becomes the Helmholtz equation for the pseudo-pressure $f(\mathbf{r})$ [25,26]

$$\nabla^2 f(\mathbf{r}) + k_0^2 f(\mathbf{r}) = -\rho_n^{-1/2}(\mathbf{r})S(\mathbf{r}) - O(\mathbf{r})f(\mathbf{r}), \tag{3}$$

where the object contrast is

$$O(\mathbf{r}) = k^2(\mathbf{r}) - k_0^2 - \rho_n^{1/2}(\mathbf{r})\nabla^2\rho_n^{-1/2}(\mathbf{r}). \tag{4}$$

The pseudo-pressure may be split into incident and scattered fields

$$f(\mathbf{r}) = f^i(\mathbf{r}) + f^s(\mathbf{r}), \tag{5}$$

where the incident field satisfies the contrast-free Helmholtz equation

$$\left[\nabla^2 + k_0^2\right]f(r) = -\rho_n^{-1/2}(\mathbf{r})S(\mathbf{r}). \tag{6}$$

If the source term S is nonzero only where $\rho_n = 1$, the incident pseudo-pressure f^i is equivalent to the true incident pressure. The homogeneous Green's function

$$g_0(\mathbf{r},\mathbf{r}') = \frac{e^{ik_0|\mathbf{r}-\mathbf{r}'|}}{4\pi|\mathbf{r}-\mathbf{r}'|} \tag{7}$$

satisfies the Helmholtz Eq. (6) in the presence of a point source located at $\mathbf{r}'$:

$$\left[\nabla^2 + k_0^2\right]g_0(\mathbf{r},\mathbf{r}') = -\delta(\mathbf{r}-\mathbf{r}'). \tag{8}$$

Hence, eliminating the incident pseudo-pressure from Eq. (3), the scattered pseudo-pressure is

$$f^s(\mathbf{r}) = \int_D d\mathbf{r}' g_0(\mathbf{r},\mathbf{r}')O(\mathbf{r}')f(\mathbf{r}'), \tag{9}$$

where D is the domain over which the object contrast O is nonzero. The scattered pseudo-pressure is equal to the true scattered pressure wherever $\rho_n = 1$. Therefore, if D is free of acoustic sources and measurements are taken outside of the scatterer domain, an integral equation solver that inverts the scattering equation

$$f(\mathbf{r}) - \int_D d\mathbf{r}' g_0(\mathbf{r},\mathbf{r}')O(\mathbf{r}')f(\mathbf{r}') = f^i(\mathbf{r}) \tag{10}$$

can neglect distinctions between the pseudo-pressure f and the true pressure p since the quantities are equivalent in the regions of interest.

3. Methods

Numerical solution of the scattering Eq. (10) is accomplished by subdividing the scattering domain D into a collection $\{c_j\colon 0 \leqslant j < N\}$ of N disjoint, cubic cells and assigning to each cell c_j a constant contrast value O_j and a constant field value f_j:

$$O(\mathbf{r}) = \sum_{j=0}^{N-1} O_j\chi_j(\mathbf{r}), \tag{11a}$$

$$f(\mathbf{r}) = \sum_{j=0}^{N-1} f_j\chi_j(\mathbf{r}), \tag{11b}$$

where χ_j is an arbitrary basis function corresponding to c_j. In this formulation, χ_j represents the characteristic function of c_j. Using the Galerkin method, the discrete expansions (11) are inserted into the scattering Eq. (10) and both sides are tested with functions χ_i to yield N discrete equations that may be represented in matrix form as

$$\mathbf{f} - \overline{\mathbf{G}}\cdot\tilde{\mathbf{f}} = \mathbf{f}^i, \tag{12}$$

where the column vectors $\mathbf{f} = [f_j]$, $\tilde{\mathbf{f}} = [O_j f_j]$; the column vector $\mathbf{f}^i = [f_j^i]$ contains entries

$$f_j^i = \int_D d\mathbf{r}\chi_j(\mathbf{r})f^i(\mathbf{r}) \tag{13}$$

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

$$G_{ij} = \int_D d\mathbf{r}\chi_i(\mathbf{r}) \int_D d\mathbf{r}' g_0(\mathbf{r},\mathbf{r}')\chi_j(\mathbf{r}'). \tag{14}$$

For large problems, iterative solution methods such as the generalized minimal residual method (GMRES) [27] or the stabilized biconjugate gradient method (BiCG-STAB) [28] are preferable to computationally expensive, direct inversion of the scattering matrix represented in (12). More importantly, iterative methods obtain solutions through repeated application of a discrete linear operator to search vectors within Krylov subspaces of the operator. With efficient methods, these operations may be performed without explicitly filling a matrix or directly computing a matrix–vector product, reducing computation time and storage demands. The fast multipole method, for example, evaluates the linear operation of Eq. (12) in $O(N)$ time and with $O(N)$ storage for N volume scattering elements [13].

3.1. The fast multipole method

A detailed discussion of the fast multipole method (FMM) and its implementation may be found in Ref. [13]. In essence, the FMM is a rank-reducing technique. Interactions between sufficiently separated groups of scattering elements are approximately represented by band-limited multipole expansions. In the Fourier domain, the multipole expansions become plane-wave expansions that are related by diagonal translation operators. An efficient solution to (12) is obtained by recursively grouping neighboring scattering elements and approximating the Green's matrix elements (14) as

$$G_{ij} \approx \int_\Omega d\hat{s} R_i(\mathbf{c}_I,\hat{s})\alpha_{I,J}(\hat{s})F_j(\mathbf{c}_J,\hat{s}), \tag{15}$$

where Ω is the unit sphere in three dimensions, R_i is the incoming far-field signature of χ_i relative to the target group center $\mathbf{c}_I$, F_j is the outgoing far-field signature (or *radiation pattern*) of χ_j relative to the source group center $\mathbf{c}_J$, and $\alpha_{I,J}$ is a diagonal translation from the source group centered at $\mathbf{c}_J$ to the target group centered at $\mathbf{c}_I$. The translator for the three-dimensional scalar Green's function (7) takes the form

$$\alpha_{I,J}(\hat{s}) = \frac{i}{4\pi}\sum_{m=0}^{L} i^m(2m+1)P_m(\hat{s}\cdot\hat{s}_{IJ})h_m(k_0|\mathbf{c}_I-\mathbf{c}_J|), \tag{16}$$

where P_m is the Legendre polynomial of order m, h_m is the spherical Hankel function of the first kind and of order m, the direction $\hat{s}_{IJ}$ is given by

$$\hat{s}_{IJ} = \frac{\mathbf{c}_I-\mathbf{c}_J}{|\mathbf{c}_I-\mathbf{c}_J|}, \tag{17}$$

and L is an appropriately selected truncation point that depends on the acoustic size of the source and target groups [13]. The approximation (15) is valid provided the source group J and the target group I are well separated, i.e., $|\mathbf{c}_I-\mathbf{c}_J|$ is larger than some multiple of the acoustic size of the source or target groups. In practice, the multiple is often chosen to be at least two, although a multiple approaching unity can be used in restricted cases [29].

The efficiency of the FMM stems from the distributive property of the translation operator. For all cells c_j within a single source group J and any target cell c_i in a group I that is well separated from J,

$$\sum_{j\in \mathrm{idx}J} G_{ij}O_j f_j \approx \int_\Omega d\hat{s} R_i(\mathbf{c}_I,\hat{s})\alpha_{I,J}(\hat{s})\sum_{j\in \mathrm{idx}J} O_j f_j F_j(\mathbf{c}_J,\hat{s}), \tag{18}$$

where idx$J$ is the collection of indices j corresponding to the cells c_j in group J. The FMM translation operation may be applied to the aggregated outgoing far-field signature $\sum_j O_j f_j F_j$ for the group, rather than individually for each source cell. Through a hierarchical grouping of the scattering elements, outgoing far-field signatures may be recursively aggregated for successively larger groups, thereby reducing the number of translation operations while increasing the number of angular samples $\hat{s}$ required to accurately represent the signatures. The translated signatures are then recursively distributed to successively smaller target groups before the incoming signatures R_i are applied at the finest level to "focus" the signatures on each individual cell. In a volume-filling scattering domain, the increased cost of aggregation, translation, and distribution for larger groups is balanced by the reduction in the number of groups for which translations must be performed, resulting in an overall scaling of $O(N)$ for an FMM involving N scattering cells [13].

For the described scattering problem, the outgoing and incoming far-field signatures of a basis or testing function χ_j are, respectively, [13]

$$F_j(\mathbf{c}_J,\hat{s}) = k_0\int_D d\mathbf{r} e^{ik_0\hat{s}\cdot(\mathbf{c}_J-\mathbf{r})}\chi_j(\mathbf{r}), \tag{19a}$$

$$R_j(\mathbf{c}_J,\hat{s}) = \int_D d\mathbf{r} e^{ik_0\hat{s}\cdot(\mathbf{r}-\mathbf{c}_J)}\chi_j(\mathbf{r}). \tag{19b}$$

Noteworthy is that, when χ_j is real-valued, the outgoing far-field signature is related to the incoming far-field signature by $F_j = kR_j^*$, where $(\cdot)^*$ denotes the complex conjugate. Hence, only one signature must be computed for χ_j. This is also the moti-

vation for requiring the field values f_j to be multiplied by O_j before multiplication with the Green's matrix in (12), rather than explicitly including the contrast O in the definition of the Green's matrix elements (14); inclusion of O in the Green's matrix would require the integrand in the far-field signature (19a) to include the factor O, breaking the symmetry of the patterns. When the discrete cells c_i are defined on a regular grid, this symmetry can be further exploited to reduce the computational and memory overhead associated with the fast multipole method. This is discussed in detail following the presentation of numerical results.

3.2. Convolution with fast fourier transforms

The fast multipole method is incapable of providing accurate approximations to the interaction between elements of neighboring (i.e., not well separated) groups. Therefore, direct evaluation of the Green's matrix elements corresponding to these interactions is required. In general formulations, interactions between neighboring groups, each containing $O(M)$ elements, are represented in local Green's matrices containing $O(M^2)$ elements each. Multiplication of these local matrices and the local portion of the vector $\tilde{\mathbf{f}}$ requires $O(M^2)$ computer time. Since M is often very small compared to the total number of scattering elements, this does not impact the complexity of the FMM. However, evaluation of neighboring interactions can still represent a significant portion of the computational effort of the FMM. When the scattering elements are cubic cells defined on a regular grid, direct evaluations can be computed in $O(M\log M)$ time, with $O(M)$ storage, using convolution via fast Fourier transforms (FFTs). If the total number of scattering elements is N, then the number of finest-level FMM groups will be $O(N/M)$. The total cost for the evaluation of neighboring interactions using FFT convolution is therefore $O(N\log M)$, in comparison to the $O(NM)$ total cost for dense-matrix multiplication of neighboring interaction matrices. Thus, for a fixed problem size N, FFT convolution dramatically reduces the dependence of calculations of neighboring interactions on the group size M.

Traditional derivations of Green's function convolution via FFT assume that interactions must be represented among all cells in the grid [20,21]. However, as illustrated in Fig. 1, FFT convolution cannot be used in this fashion to compute neighboring interactions in the FMM with immediately adjacent groups. A convolutional grid that encompasses only the neighbors of Target A will compute interactions at Target B that do not include the influence of its exclusive neighbors (the white boxes). Defining the convolutional grid to encompass all illustrated boxes results in the exclusive neighbors of Target A (the light-gray boxes) inappropriately influencing Target B, as well as the exclusive neighbors of Target B inappropriately influencing Target A. Furthermore, neither of these grids properly accommodate the unlabeled boxes in the diagram when those boxes have neighbors not shown in the illustration.

Modifications of the FFT convolution scheme for pairwise group interactions are here derived in the context of acoustic interactions between finest-level FMM groups. Let c_i represent a target cell in the FMM group I, while J represents a collection of source cells in the FMM group (the groups I and J may be the same). Furthermore, let each group contain M scattering cells, arranged in an $M_x \times M_y \times M_z$ grid. The cell c_i will have an index representation (l,m,n) within the grid local to group I, while a cell c_j will have an index representation (t,u,v) within the grid local to group J. The scattering contribution (9) to the matrix–vector product (12) at c_i due to sources within group J may be expressed as

$$f_i^s = \sum_{j\in \mathrm{idx}J} G_{ij}O_j f_j = \sum_{t=0}^{M_x-1}\sum_{u=0}^{M_y-1}\sum_{v=0}^{M_z-1} G_{IJ}(l-t,m-u,n-v)O_{tuv}f_{tuv}, \tag{20}$$

where $O_{tuv}f_{tuv} = O_j f_j$ for the global index j corresponding to the local grid index (t,u,v) and the pairwise Green's function

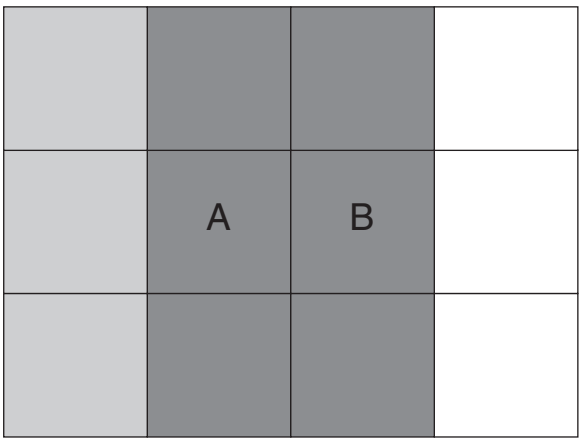

Fig. 1. Neighbor interactions for target boxes A and B. Dark-gray neighbors influence both targets; light-gray neighbors influence only target A; white neighbors influence only target B.

$$G_{IJ}(l-t,m-u,n-v)=\int_D d\mathbf{r}\,\chi_0(\mathbf{r})\int_D d\mathbf{r}'\,g_0(\mathbf{r},\mathbf{r}'+\Delta\mathbf{r}+\mathbf{r}_{JI})\chi_0(\mathbf{r}). \tag{21}$$

Translational invariance and the equivalence of all cells c_i allows the characteristic functions χ_i and χ_j in the definition of the Green's matrix element G_{ij} (14) to be replaced with an arbitrary characteristic function such as χ_0. Hence, elements of the pairwise Green's function G_{IJ} are functions only of the separation of two cells, rather than their absolute positions. The offsets in (21) are given by

$$\mathbf{r}_{JI}=\mathbf{c}_J-\mathbf{c}_I, \tag{22a}$$
$$\Delta\mathbf{r}=[(l-t)\Delta x,(m-u)\Delta y,(n-v)\Delta z], \tag{22b}$$

where each scattering cell has x, y, and z lengths Δx, Δy, and Δz, respectively.

The three-dimensional convolution (21) can only be evaluated in the Fourier domain if G_{IJ} satisfies the cyclic properties

$$G_{IJ}(-t,u,v)=G_{IJ}(M_x-t,u,v), \tag{23a}$$
$$G_{IJ}(t,-u,v)=G_{IJ}(t,M_y-u,v), \tag{23b}$$
$$G_{IJ}(t,u,-v)=G_{IJ}(t,u,M_z-v). \tag{23c}$$

To satisfy these criteria, the pairwise Green's function G_{IJ} defined on the local $M_x\times M_y\times M_z$ grid must be replaced with a modified Green's function G'_{IJ} defined on an expanded $2M_x\times 2M_y\times 2M_z$ grid. The modified Green's function

$$G'_{IJ}(t,u,v)=G_{IJ}(t',u',v'), \tag{24}$$

in which the mappings

$$t'=\begin{cases}t & \text{if } 0\leqslant t<M_x,\\ t-2M_x & \text{if } M_x\leqslant t<2M_x,\end{cases} \tag{25a}$$
$$u'=\begin{cases}u & \text{if } 0\leqslant u<M_y,\\ u-2M_y & \text{if } M_y\leqslant u<2M_y,\end{cases} \tag{25b}$$
$$v'=\begin{cases}v & \text{if } 0\leqslant v<M_z,\\ v-2M_z & \text{if } M_y\leqslant v<2M_z\end{cases} \tag{25c}$$

relate indices (t',u',v') on the original $M_x\times M_y\times M_z$ grid to indices (t,u,v) on the expanded $2M_x\times 2M_y\times 2M_z$ grid. Note that the index mappings (25) are different than those used in traditional FFT convolution formulations because the pairwise Green's function is not symmetric about zero. However, when groups I and J coincide (such that $r_{JI}=0$) the two expansions are equivalent. The product of the contrast and pseudo-pressure must similarly be expressed on an expanded grid as

$$O'_{tuv}f'_{tuv}=\begin{cases}O_{tuv}f_{tuv}, & \text{if } 0\leqslant t<M_x,\ 0\leqslant u<M_y,\ \text{and } 0\leqslant v<M_z;\\ 0 & \text{otherwise.}\end{cases} \tag{26}$$

The expressions (24) and (26) allow the contribution of group J to c_i (20) to be represented as a Fourier-domain multiplication:

$$f_i^s=\mathbb{F}^{-1}\left[\mathbb{F}(O'_{tuv}f'_{tuv})\cdot\mathbb{F}G'_{IJ}(t,u,v)\right]_{lmn}, \tag{27}$$

where $\mathbb{F}$ represents the discrete Fourier transform and the group-local index triple (l,m,n) corresponds to the global index i. These contributions can be summed for all neighbor groups J of I to calculate the near terms overlooked by the fast multipole method.

4. Numerical results

Numerical experiments were performed on a single scattering geometry under a variety of conditions. The geometry consisted of twelve scattering spheres embedded in an infinite, lossless background (water) with a sound speed of 1509 m/s and no attenuation. Each of the spheres had a unique, uniform material profile designed to mimic the human tissues listed in Table 1. The properties of each of the twelve spheres are listed in Table 2, and a diagram depicting the sphere arrangement is shown in Fig. 2.

Table 1
Material properties of spheres designed to mimic human tissue.

Tissue	Sound speed m/s	Absorption dB/(cm MHz)
Water	1509.0	0.00
Fat	1478.0	0.52
Muscle	1547.0	0.91
Skin	1613.0	1.61

Table 2
Characteristics of the spheres in the tissue-mimicking phantom.

Radius (mm)	Center (mm)			Tissue
	x	y	z	
4.0	0.0	0.0	0.0	Fat
5.0	14.0	2.0	4.0	Skin
5.0	5.0	−10.0	−4.0	Fat
3.0	17.0	−7.0	0.0	Fat
7.8	−10.0	10.0	7.2	Muscle
7.8	5.0	12.0	−7.2	Muscle
5.0	14.0	12.0	3.0	Muscle
5.0	−5.0	−18.0	−3.0	Skin
2.5	7.5	−18.0	−2.0	Skin
1.5	−4.0	20.0	0.0	Skin
2.5	−18.0	4.0	2.0	Skin
9.1	−12.5	−5.0	−5.2	Muscle

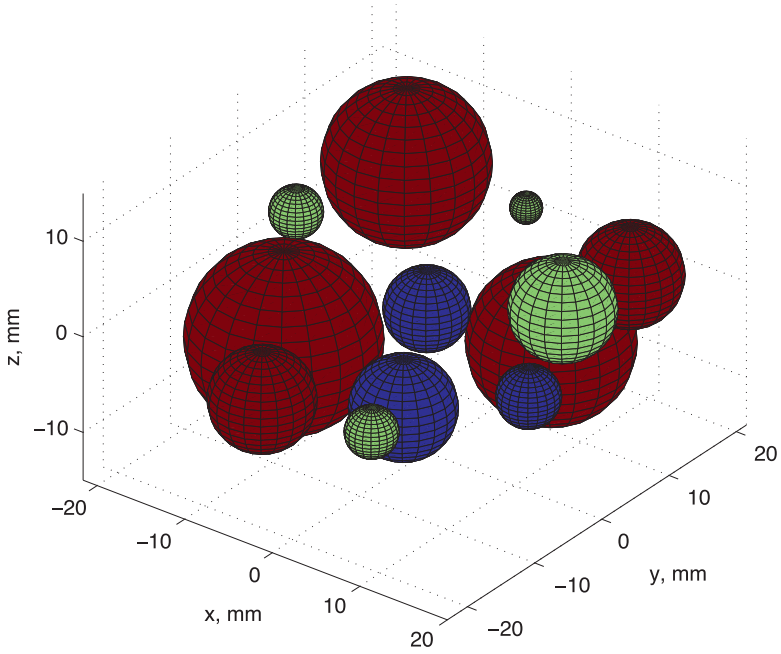

Fig. 2. An arrangement of twelve tissue-mimicking spheres in a water background. The spheres are colored according to their composition as listed in Table 2. Red, green and blue represent skin, muscle, and fat, respectively. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

The collection of scattering spheres was chosen because it is easy to model and because efficient methods can be developed [29,30] to compute solutions that may be used to validate the FMM. Because the FMM is designed to compute scattering from domains with arbitrarily variable medium characteristics, it can not exploit efficiencies associated with this particular geometry and is, therefore, not the optimum choice for these computations. However, the general nature of the FMM means that the performance observed in these experiments is representative of the performance that would be expected from scattering computations involving truly arbitrary geometries, for which the FMM would be a more suitable choice.

The twelve-sphere arrangement was used in two numerical experiments. In the first, the computational scaling of the FMM was examined using both dense-matrix multiplication and FFT convolution for neighboring interactions as a function of increasing finest-level group size. In the second experiment, a higher incident frequency was used to produce a problem of much larger scale. This problem was solved using both the FMM with FFT convolution and compared to the results computed with a fast sphere solver [29]. In both experiments, iterative inversion of the scattering equation was performed using the stabilized biconjugate gradient method (BiCG-STAB) [28]. The BiCG-STAB solver requires two matrix–vector products and four inner products per iteration.

4.1. Computational scaling as a function of Finest-Level group size

Investigation of FMM scaling with FFT convolution was accomplished using a relatively small-scale problem to facilitate rapid solution of multiple scattering experiments involving finest-level groups containing variable numbers of scattering elements. Scattering from the twelve-sphere arrangement was simulated at 300 kHz, corresponding to a water wavelength $\lambda_0 = 5.03$ mm. The scattering domain was confined to a cube with sides of length $9.6\lambda_0$ and was sampled with a regular spacing of $0.1\lambda_0$ per dimension, resulting in a $96 \times 96 \times 96$ grid of scattering elements. The water background accounts for approximately 90% of the total volume of the scattering domain. In practical applications, the FMM can be exploited to avoid meshing this large volume of non-scattering material and, therefore, to improve performance by a factor of ten. However, for this experiment, the FMM implementation was kept uncomplicated by meshing the entire volume. While a pure CG-FFT implementation might be more suitable for this non-optimal grid, the FMM performance observed is representative of the performance of problems involving the same number of elements distributed throughout disjoint grids for which the CG-FFT is not a suitable choice.

The sizes of the finest-level FMM groups were constrained to be integer multiples of the unknown grid spacing in all three dimensions. Thus, for every test configuration, unknowns in each finest-level group coincided with the same affine grid. This was required to permit FFT convolution of neighboring interactions and also reduced the memory and setup demands of the FMM. Groups were chosen to be cubes containing a number $2 \leqslant M_d \leqslant 16$ of scattering elements per dimension. The total number of elements in a finest-level group is denoted $M = M_d^3$.

Timing experiments were run on a computer system housing 16 GB of RAM and two dual-core processors operating at 2.4 GHz. Implementations of the FMM with FFT convolution and dense-matrix multiplication for neighboring interactions were each multithreaded to take advantage of all four processing cores with 97% parallel efficiency. For both implementations of the FMM, times were recorded for computation of a full matrix–vector product, the evaluation of only neighboring interactions, and the computation of far-field signatures at the finest level. Far-field signature computation at the finest level consists of evaluation of outgoing far-field signatures for groups of scattering elements, as well as distribution of incoming signatures for each group to its constituent elements. Only the time spent computing neighboring interactions varied between FMM implementations; all other components of the FMM computation were identical. The recorded computation times are shown in Fig. 3 as a function of the number of scattering elements in each finest-level FMM group. In addition, the proportion of the total matrix–vector-product time spent computing neighboring interactions and far-field signatures is shown in Fig. 4 for each FMM implementation.

For $M_d \leqslant 4$, there was no appreciable difference in computation time for a matrix–vector product whether dense-matrix multiplication or FFT convolution was used to represent neighboring interactions. However, as M_d increased, dense-matrix multiplication of neighboring interactions became the dominant portion of the matrix–vector product, rapidly increasing the overall computation time. The FMM implementation employing FFT convolution for neighboring interactions was more efficient and exhibited a reduced dependence on the number of scattering elements in each finest-level FMM group.

Over the range $4 \leqslant M_d \leqslant 12$ (corresponding to a 27-fold increase in the total number of elements functions per FMM group), the total time required to compute a matrix–vector product using FFT convolution for neighboring interactions did not significantly increase. Fluctuations over this range of group sizes were attributable to the changing FMM hierarchy and to differences in the efficiency of computing FFTs of different sizes (*e.g.*, FFTs with lengths that are powers of two are more efficient than FFTs with lengths that are products of powers of other prime numbers).

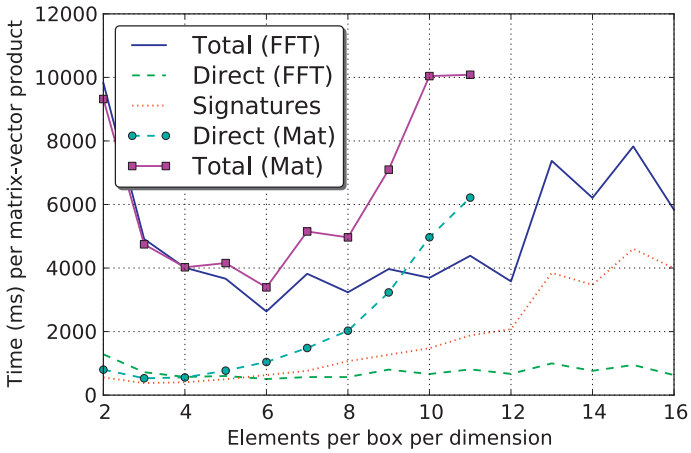

Fig. 3. Time required to compute a complete matrix–vector product ("Total") as a function of finest-level FMM group size for both FFT convolution ("FFT") and dense-matrix multiplication ("Mat") of neighboring interactions. Also shown are the times for computing only neighboring interactions ("Direct") and for computing only far-field signatures ("Signatures") at the finest level.

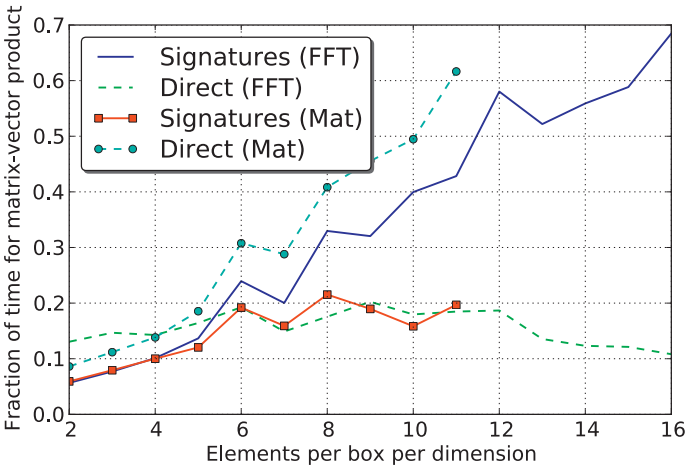

Fig. 4. Fractions of total time for matrix–vector products consumed by far-field signatures ("Signatures") and neighboring interactions ("Direct") when using the FMM with FFT convolution ("FFT") and dense-matrix multiplication ("Mat").

For $M_d > 12$, the cost of far-field signature computations at the finest level became the dominant portion of the matrix–vector product when FFT convolutions are used to compute neighboring interactions. The number of samples required to represent the far-field signatures for a single FMM scales as $O(M_d^2)$. As Figs. 3 and 4 show, the efficiency of the hierarchical FMM suffers long before the efficiency of FFT convolution for neighboring interactions as the finest-level group size is increased. Hence, for optimum efficiency, the FMM should not be used only for interactions between large, dense scatterers; these scatterers should themselves be subdivided.

4.2. Validation for a larger problem

To produce a larger problem for validation of the FMM, an incident wave of 1 MHz was assumed, corresponding to a water wavelength $\lambda_0 = 1.51$ mm. At this frequency, the scattering geometry was confined to a cube with sides of length $32\lambda_0$. Sampled with a regular spacing of $0.1\lambda_0$ per dimension, the resulting $320 \times 320 \times 320$ grid consists of almost 33 million unknowns.

By taking advantage of the regular structure of unknowns and FFT convolution, FMM setup time (computing the diagonal translators for every level, interpolation matrices to resample far-field signatures, far-field signatures for finest-level FMM groups, and Green's functions for neighboring interactions at the finest level) was reduced to less than 1 min. Total memory consumption was limited to just under 6 GB, approximately 2 GB of which was reserved to store the Fourier transforms of pressure distributions in all boxes to avoid repeated computation. A more aggressive optimization would dramatically reduce this cache size without a significant (or any) increase in computation time by carefully arranging neighboring calculations so that only currently-required Fourier transforms were stored, an were fully utilized before being discarded.

Computation of scattering from a single plane wave incident from the $+x$ direction required 14 BiCG-STAB iterations to converge to a residual norm of 10^{-3}, and was completed in 5,716 s. The finest-level FMM groups contained 5 basis functions per dimension, resulting in a six-level FMM in which all finest-level groups are uniformly populated. The resulting azimuthal scattering pattern is shown in Fig. 5. The computed pattern was compared to the fast sphere scattering solver presented in Ref. [29] by computing the overall root-mean-squared error (RMSE), given by

$$\mathrm{RMSE} = \left[\int_\Omega d\hat{s} \left| p_s(\hat{s}) - p_{s,r}(\hat{s}) \right|^2 \Big/ \int_\Omega d\hat{s} \left| p_{s,r}(\hat{s}) \right|^2 \right]^{1/2} \tag{28}$$

for a far-field scattered pressure p_s and a reference scattered pressure $p_{s,r}$ on the unit sphere Ω. The RMSE between the two methods was 2.7%. The fast sphere solver, by exploiting this simplistic scattering geometry, is much more efficient for this particular problem. The reference solution was computed in only 13 s.

5. Discussion

Two aspects of the FMM efficiency merit special comment. The first is the reduction in storage and setup time that is gained by using a regular arrangement of scattering elements and the use of FFT convolution for neighboring interactions. The second is the potential to further improve computational efficiency of the matrix–vector product when the finest-level FMM group size is large.

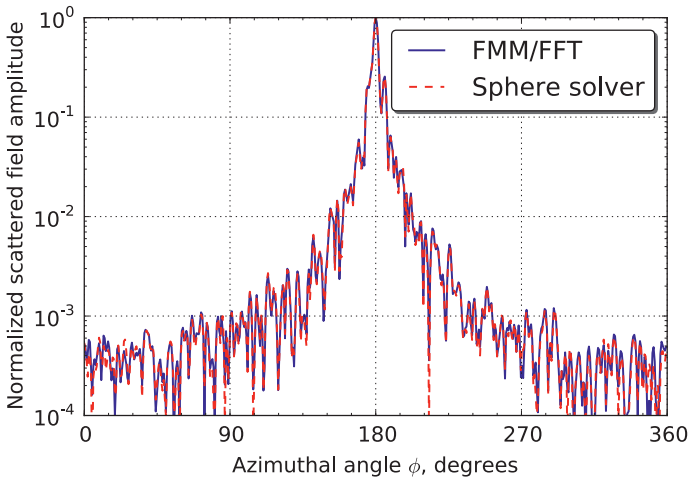

Fig. 5. Azimuthal scattering pattern of the twelve-sphere, tissue-mimicking phantom for a 1 MHz plane wave incident from the $+x$ direction.

5.1. Storage and setup efficiency

As noted in the validation experiment for a 1 MHz incident plane wave, the FMM setup time was reduced to less than 1 min by taking advantage of the regular arrangement of basis functions in the volume scattering problem. In typical applications of the FMM, basis functions of arbitrary shapes and locations in either two or three dimensions yield incoming and outgoing far-field signatures that may be unique for every basis function. To reduce the computational cost of applying the matrix–vector product, these patterns are usually precomputed and stored for every basis function. For the example above, this requires evaluation of nearly 66 million far-field signatures, with 244 samples per pattern. However, since the incoming signature is the complex conjugate of the outgoing signature, only 33 million signatures required computation.

When the unknowns are defined on a regular grid and made to coincide with the grid defined by the finest-level FMM groups, the far-field signatures need only be computed for a single, representative box. In the experiment above, each box contains only $5^3 = 125$ scattering elements. Hence, the cost of computing far-field signatures, normally one of the two costliest parts of FMM setup, is reduced by a factor of 262,144. The memory is also reduced by the same factor; the cost for storing the representative signatures in floating-point arithmetic is only 238 kB, whereas the cost for storing all unique patterns would be nearly 60 GB!

Neighboring interactions, the other of the two costliest parts of FMM setup, also need be computed for only a representative set of interactions. Because not every finest-level group is bordered on all sides by other groups, the reduction in memory and computation time is not exactly N/M for a problem involving N unknowns and M unknowns in each finest-level group. However, the reduction is still $O(N/M)$, with a proportionality constant close to unity.

5.2. Further improving efficiency for large groups

The results in Figs. 3 and 4 suggest that modifications to the FMM can improve efficiency when the group size is large. As the group size is increased, the cost of evaluating far-field signatures exceeds the cost of evaluating neighboring interactions. However, the evaluation of far-field signatures and neighboring interactions can be performed at different levels in the FMM hierarchy. If group sizes are made large to avoid low-frequency breakdown of outgoing-to-incoming diagonal translation, they may be further subdivided to compute outgoing far-field signatures or to use incoming signatures to distribute fields to basis functions, because these operations are not subject to the same low-frequency breakdown problems.

If L represents the number of levels in an FMM hierarchy, and neighboring interactions are evaluated at some level $L_N < L$, the computational cost of all FMM operations (translations and resampling between levels) is strictly bounded above by the cost of FMM operations in an ordinary L-level FMM because outgoing-to-incoming translations will not be performed for levels greater than L_N. Thus, the total cost for performing a matrix–vector product in this fashion does not exceed the sum of the cost of FMM operations for an ordinary L-level FMM and the cost of computing neighboring interactions in an L_N-level FMM.

As an example, consider the scaling experiment in Fig. 3. For the FMM structure with 8 elements per group per dimension, the total time required to compute a matrix–vector product on the aforementioned four-core, parallel system was 71.2 s, of which 12.5 s was dedicated to neighboring interactions with FFT convolution. Hence, 58.7 s was required to compute all FMM related interactions. If the neighboring interactions are instead evaluated on the next-highest FMM level, with 16 elements per group per dimension, the time for FMM computations will not exceed 58.7 s, and the cost of neighboring interactions will increase to 13.9 s as measured in the experiment. The total time for the matrix–vector product in this setup is

bounded from above by 72.6 s, which is significantly lower than the 128.3 s observed when the finest FMM level enclosed 16 elements per group per dimension. Thus, evaluation of neighboring interactions by using FFT convolution at a coarser FMM level can simultaneously avoid the increasing cost of far-field signature evaluations for increasing finest-level box size and the low-frequency breakdown problem when FMM groups are acoustically small.

6. Conclusion

An integral equation formulation was presented for the solution of acoustic scattering problem in the presence of three-dimensional, penetrable scatterers categorized by arbitrary spatial variations of sound speed, attenuation, and density. The arrangement of scattering elements in a regular grid and the use of the FMM facilitates a dramatic reduction in storage requirements and the computational cost for the FMM setup. Additionally, the regular arrangement of elements permit the use of FFT convolution for evaluation of neighboring interactions, reducing the penalty incurred when finest-level FMM groups are relatively large. Modification of the method to compute direct interactions on a coarser FMM level offers a further-reduced dependency of computation time on the finest-level group structure.

Acknowledgments

Professor Weng C. Chew, of the University of Illinois and the University of Hong Kong, is thanked for the use of his `ScaleME` parallel FMM library [13,31]. Jason C. Tillett and Jeffrey P. Astheimer, both of the Ultrasound Research Laboratory at the University of Rochester, are thanked for helpful discussions, suggestions, and comments about material in this paper. This research was funded in part by NIH Grants EB 00280 and EB 009692 and the University of Rochester Diagnostic Ultrasound Research Laboratory Industrial Associates.

References

[1] M. Tabei, T.D. Mast, R.C. Waag, Simulation of ultrasonic focus aberration and correction through human tissue, J. Acoust. Soc. Am. 113 (2003) 1166–1176.
[2] J.C. Lacefield, W.C. Pilkington, R.C. Waag, Distributed aberrators for emulation of pulse distortion by abdominal wall, Acoust. Res. Lett. Online 3 (2002) 47–52.
[3] T.D. Mast, L.M. Hinkelman, L.A. Metlay, M.J. Orr, R.C. Waag, Simulation of ultrasonic pulse propagation, distortion, and attenuation in the human chest wall, J. Acoust. Soc. Am. 106 (1999) 3665–3677.
[4] A.J. Devaney, M.L. Oristaglio, Inversion procedure for inverse scattering within the distorted-wave Born approximation, Phys. Rev. Lett. 51 (1983) 237–240.
[5] W.C. Chew, Y.M. Wang, Reconstruction of two-dimensional permittivity distribution using the distorted Born iterative method, IEEE Trans. Med. Imag. 9 (1990) 218–225.
[6] R. Lavarello, M. Oelze, Density imaging using inverse scattering, J. Acoust. Soc. Am. 125 (2009) 793–802.
[7] A.J. Hesford, W.C. Chew, Fast inverse scattering solutions using the distorted Born iterative method and the multilevel fast multipole algorithm, J. Acoust. Soc. Am. 128 (2) (2010) 679–690.
[8] T.D. Mast, A.I. Nachman, R.C. Waag, Focussing and imaging using eigenfunctions of the scattering operator, J. Acoust. Soc. Am. 102 (1997) 715–726.
[9] F. Lin, A.I. Nachman, R.C. Waag, Quantitative imaging using a time-domain eigenfunction method, J. Acoust. Soc. Am. 108 (2000) 899–912.
[10] R.C. Waag, F. Lin, T.K. Varslot, J.P. Astheimer, An eigenfunction method for reconstruction of large-scale and high-contrast objects, IEEE Trans. Ultrason. Ferroelectr. Freq. Control 54 (2007) 1316–1332.
[11] L. Greengard, V. Rokhlin, A fast algorithm for particle simulations, J. Comput. Phys. 73 (1987) 325–348.
[12] V. Rokhlin, Rapid solution of integral equations of scattering theory in two dimensions, J. Comput. Phys. 86 (1990) 414–439.
[13] W.C. Chew, J. Jin, E. Michielssen, J. Song (Eds.), Fast and Efficient Algorithms in Computational Electromagnetics, Artech House, Boston, 2001.
[14] E. Michielssen, J.-M. Jin, Guest editorial for the special issue on large and multiscale computational electromagnetics, IEEE Trans. Antennas Propag. 56 (8) (2008) 2146–2149.
[15] M. Bebendorf, Approximation of boundary element matrices, Numer. Math. 86 (2000) 565–589.
[16] K. Zhao, M.N. Vouvakis, J.-F. Lee, The adaptive cross approximation algorithm for accelerated method of moments computations of EMC problems, IEEE Trans. Electromagn. Compat. 47 (2005) 763–773.
[17] J. Shaeffer, Direct solve of electrically large integral equations for problem sizes to 1M unknowns, IEEE Trans. Antennas Propag. 56 (8) (2008) 2306–2313.
[18] H. Zhai, Q. Chen, Q. Yuan, K. Sawaya, C. Liang, Analysis of large-scale periodic array antennas by CG-FFT combined with equivalent sub-array preconditioner, IEICE Trans. Commun. 1 (2006) 922–928.
[19] T.J. Cui, W.C. Chew, A.A. Aydiner, Y.H. Zhang, Fast-forward solvers for the low-frequency detection of buried dielectric objects, IEEE Trans. Geosci. Remote Sens. 41 (2003) 2026–2036.
[20] X.M. Xu, Q.H. Liu, The BCGS-FFT method for electromagnetic scattering from inhomogeneous objects in a planarly layered medium, IEEE Antennas Wireless Propag. Lett. 1 (2002) 77–80.
[21] T.J. Cui, W.C. Chew, Fast algorithm for electromagnetic scattering by buried 3-D dielectric objects of large size, IEEE Trans. Geosci. Remote Sens. 37 (1999) 2597–2608.
[22] J. De Zaeytijd, I. Bogaert, A. Franchois, An efficient hybrid MLFMA-FFT solver for the volume integral equation in case of sparse 3D inhomogeneous dielectric scatterers, J. Comput. Phys. 227 (2008) 7052–7068.
[23] F.A. Duck (Ed.), Physical Properties of Tissue: A Comprehensive Reference Book, Academic Press, London, 1990.
[24] P.M. Morse, K.U. Ingard, Theoretical Acoustics, McGraw-Hill, New York, 1968.
[25] S.A. Johnson, F. Stenger, C. Wilcox, J. Ball, M.J. Berggren, Wave equations and inverse solutions for soft tissue, Acoust. Imaging 11 (1982) 409–424.
[26] S. Pourjavid, O.J. Tretiak, Numerical solution of the direct scattering problem through the transformed wave equation, J. Acoust. Soc. Am. 91 (1992) 639–645.
[27] Y. Saad, M.H. Schultz, GMRES—a generalized minimal residual algorithm for solving nonsymmetric linear-systems, SIAM J. Sci. Stat. Comput. 7 (1986) 856–869.
[28] H.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 (2) (1992) 631–644.

[29] A.J. Hesford, J.P. Astheimer, L.F. Greengard, R.C. Waag, A mesh-free approach to acoustic scattering from multiple spheres nested inside a large sphere by using diagonal translation operators, J. Acoust. Soc. Am. 127 (2) (2010) 850–861.
[30] A.J. Hesford, J.P. Astheimer, R.C. Waag, Acoustic scattering by arbitrary distributions of disjoint, homogeneous cylinders or spheres, J. Acoust. Soc. Am. 127 (5) (2010) 2883–2893.
[31] S. Velamparambil, W.C. Chew, Analysis and performance of a distributed memory multilevel fast multipole algorithm, IEEE Trans. Antennas Propag. 53 (2005) 2719–2727.