A fast boundary element method for the scattering analysis of high-intensity focused ultrasound

Elwin van 't Wout[a)]
Department of Computer Science, University College London, London, United Kingdom

Pierre Gélat
Department of Mechanical Engineering, University College London, London, United Kingdom

Timo Betcke
Department of Mathematics, University College London, London, United Kingdom

Simon Arridge
Department of Computer Science, University College London, London, United Kingdom

High-intensity focused ultrasound (HIFU) techniques are promising modalities for the non-invasive treatment of cancer. For HIFU therapies of, e.g., liver cancer, one of the main challenges is the accurate focusing of the acoustic field inside a ribcage. Computational methods can play an important role in the patient-specific planning of these transcostal HIFU treatments. This requires the accurate modeling of acoustic scattering at ribcages. The use of a boundary element method (BEM) is an effective approach for this purpose because only the boundaries of the ribs have to be discretized instead of the standard approach to model the entire volume around the ribcage. This paper combines fast algorithms that improve the efficiency of BEM specifically for the high-frequency range necessary for transcostal HIFU applications. That is, a Galerkin discretized Burton–Miller formulation is used in combination with preconditioning and matrix compression techniques. In particular, quick convergence is achieved with the operator preconditioner that has been designed with on-surface radiation conditions for the high-frequency approximation of the Neumann-to-Dirichlet map. Realistic computations of acoustic scattering at 1 MHz on a human ribcage model demonstrate the effectiveness of this dedicated BEM algorithm for HIFU scattering analysis.
© 2015 Acoustical Society of America. [http://dx.doi.org/10.1121/1.4932166]

I. INTRODUCTION

The liver is a common site of occurrence for tumors,[1,2] and the incidence of liver cancer is on the rise in Europe.[3] Hepatocellular carcinoma, the most common form of liver cancer, is a growing global health problem as it is the third most common cause of cancer related death[4] with increased incidence rates worldwide.[5]

The first choice of therapy for liver cancer is either surgical resection or transplantation.[2,6] The suitability of a patient for surgical resection is highly dependent on the size, location, and number of tumors.[2,7] The prognosis for patients having undergone a resection remains poor, often due to the fact that other tumors may have been present during surgery but remained undetected due to their small size.[2,8] Moreover, the risks associated with conventional surgical treatments render them unsuitable for the majority of patients: resection is an invasive procedure that involves the loss of large amounts of blood. Thus the ability to ablate tumors accurately and non-invasively within the liver will have significant clinical impact.

High-intensity focused ultrasound (HIFU) is a medical procedure that uses high-amplitude ultrasound to heat and ablate a localized region of tissue. High energy may be accurately targeted within a well-defined and predetermined volume, and tissue destruction may be achieved without damaging the overlying tissue.[9] The feasibility of HIFU for the treatment of a range of different cancers, including those of the liver, has been demonstrated.[10,11] As a non-invasive focused therapy, HIFU offers significant advantages over surgical resection, and it has been shown that it may serve as a promising locally ablative technique for the treatment of hepatocellular carcinoma.[12,13] Despite this, there are a number of significant challenges that currently hinder its more widespread clinical application. For instance, rib bone both absorbs and reflects ultrasound strongly. Hence a common side effect of focusing ultrasound in regions located behind the ribcage is the overheating of bone and surrounding tissue, which can lead to skin burns.[14,15] Furthermore, the presence of ribs can lead to aberrations at the focal region due to effects of diffraction,[16] which can lead to deterioration of the focusing effect of the incident acoustic field. Aubry *et al.*[17] have stated that one of the minimal technical specifications of a HIFU system for the treatment of liver tumors should be to transmit energy in between, below, or through the ribs without damaging the ribs or causing a skin burn. This

a)Electronic mail: e.wout@ucl.ac.uk

requirement is likely to rely on a patient-specific treatment planning protocol that features numerical modeling to predict and optimize the acoustic field near a ribcage.

Different soft tissue types in the human body tend to bear similar acoustic properties. The speed of propagation of longitudinal waves is generally comparable in different soft tissues and is 1500 m/s approximately.[18] The same is true of the density,[18] which is around 1000 kg/m³. Whilst it is known that soft tissue heterogeneities can lead to aberrations of the focus,[19] a first step toward treating the problem of acoustic scattering by the ribcage is to consider the ribs to be immersed in an infinite homogeneous medium.

There have been several attempts to model HIFU fields in the presence of ribs, but full-wave three-dimensional (3D) modeling of ultrasound propagation in the presence of bone still presents substantial computational challenges. As a result of the large domain dimensions at the MHz frequencies required for transcostal HIFU, many models have relied on simplified shadowing techniques.[15,16,20,21] Whilst these algorithms, such as ray tracing, may replicate features of wave propagation during transcostal HIFU treatments, they do not accurately address the actual scattering mechanisms involved in complex 3D structures and are likely to be of limited use in clinical treatment planning applications.

Finite-difference time-domain (FDTD) schemes have been used in transcostal HIFU simulations.[22,23] Nevertheless, owing to the large grid sizes resulting from the discretization of the entire region around the ribcage, these studies have been confined to 2D models. Computationally efficient approaches, such as k-space pseudospectral methods have shown promise for modeling acoustic fields in heterogeneous media.[19,24–26] Whilst these methods can deal somewhat with weak inhomogeneities in the propagating medium, an increasingly finer computational grid is required at an interface of soft tissue and bone.[19]

Gélat et al.[27] proposed a boundary element method (BEM) approach to model the scattering of a HIFU field by human ribs. In BEM, the exterior scattering problem is reformulated into an integral equation on the surface of the object. This allows for accurate discretization of the boundary conditions at the interface and a natural representation of unbounded regions without truncation of the computational domain. The dimensional reduction, in combination with fast computational algorithms, results in efficient simulation of large-scale problems. Being based on Green's function representations, the BEM is devoid of numerical dispersion and dissipation effects. Furthermore, it is applicable for both reflecting and absorbing boundary conditions on the ribs. The BEM is therefore particularly suited to transcostal HIFU simulations. Whilst the BEM approach is linear and does not account for effects due to nonlinear propagation in tissue, it has the advantage of modeling the scattering of an acoustic field by arbitrary 3D rib topologies efficiently.

The capability of the BEM as a forward model within an optimization method for the design of a transducer array for transcostal HIFU applications has already been demonstrated.[27–29] This algorithm uses the Burton–Miller formulation, which has the advantage of being devoid of spurious resonance modes.[30] However, the efficiency of the collocation discretization deteriorates for large-scale problems, where the wavelength is small compared to the size of the object. Frequencies in the MHz range are necessary for transcostal HIFU applications; this leads to substantial requirements on computational resources in terms of both storage and time. The scattering analysis on the ribcage with the BEM has been observed to be the computational bottleneck of the optimization algorithm for the design algorithm of HIFU transducers. To alleviate the computational requirements, a dedicated fast algorithm that extends the applicability of BEM to high-frequency acoustic scattering will be proposed in this paper. Specifically, the method is based on a Galerkin discretization in combination with preconditioning and compression techniques.

A main reason for the deterioration in efficiency with increasing frequency is the differentiating property of the hypersingular operator within the Burton–Miller formulation. A regularizing preconditioner will be used that results in a second-kind Fredholm integral equation and thus rapid convergence. The preconditioner is based on a high-frequency approximation of the Neumann-to-Dirichlet map with on-surface radiation conditions (OSRC).[31,32] The discretization of this operator preconditioner results in a sparse matrix and therefore relatively little overhead during each iteration. Moreover, fast algorithms for the matrix-vector multiplication such as $\mathcal{H}$-matrix compression techniques and fast multipole methods (FMM) can straightforwardly be combined with the preconditioner.[33] In this study, the $\mathcal{H}$-matrix technique has been used to compress the storage of the matrices arising from the Galerkin discretization. This algorithm is very efficient for the frequency range of typical HIFU configurations and can, once assembled, be used efficiently for multiple right-hand-side vectors. These characteristics make it a feasible method for the future goal of optimizing HIFU transducer arrays with the BEM as a forward model. For very high frequencies though, specialized FMM techniques will probably be the preferred choice.[34] Finally, the open-source package BEM++ has been used as an implementation platform, which has $\mathcal{H}$-matrix compression functionality.[35]

This paper starts with a presentation of the model formulation in Sec. II, including the boundary integral representation and the model for HIFU fields from multi-element transducer arrays. The algorithms to improve the efficiency of the BEM will be explained in Sec. III, in particular, the OSRC preconditioning and $\mathcal{H}$-matrix compression techniques. Numerical simulations on perfectly rigid ribs immersed in water will be described in Sec. IV. This case is particularly relevant to HIFU experimental *ex vivo* studies involving ribs.[23] It is also applicable to the characterisation of HIFU fields in the presence of rib phantoms immersed in water. The experimental results in this paper demonstrate the convergence improvement of the preconditioner, the compression rates of the matrices, and its feasibility for the simulation of transcostal HIFU techniques.

II. MODEL FORMULATION

The aim of this paper is to develop an efficient computational method for simulation of transcostal HIFU modalities.

Of special interest is the accurate modeling of the acoustic scattering on ribcages for which a model of a rigid body immersed in an infinite fluid will be used. To this end, consider an object representing a part of a ribcage residing in a homogeneous lossless medium for which the physical parameters of water will be used. Although ribs are both absorbing and reflective, they are assumed to be perfectly rigid in this paper. It should be noted that the BEM naturally allows for an extension to absorbing objects. The HIFU transducers typically operate in a small frequency band in the MHz range. This allows for the use of time-harmonic waves with a fixed wavelength.

A. Exterior wave formulation

The propagation of acoustic waves in homogeneous media can be described by a system of the Helmholtz equation and boundary conditions. Let us consider a geometry with a bounded domain $\Omega_{\text{int}} \subset \mathbb{R}^3$ representing a scatterer and an exterior domain $\Omega_{\text{ext}} = \mathbb{R}^3 \backslash \overline{\Omega_{\text{int}}}$. The interior domain can be nonconvex and disconnected, but its boundary, denoted by $\Gamma = \partial\Omega_{\text{int}}$, is assumed to be Lipschitz continuous. The unit normal $\hat{\mathbf{n}}$ on Γ is outward pointing, i.e., from Ω_{int} toward Ω_{ext}. The wavenumber of the time-harmonic acoustic wave is given by $k = 2\pi/\lambda$, where λ denotes the wavelength. For a rigid object, the scattered field pressure p of an acoustic wave from an excitation p_{inc} can be modeled with the Helmholtz exterior boundary value problem

$$
\begin{cases}
\Delta p + k^2 p = 0 & \text{in } \Omega_{\text{ext}}, \\
\partial_n (p + p_{\text{inc}})|_\Gamma = 0 & \text{on } \Gamma, \\
\lim_{|\mathbf{x}|\to\infty} |\mathbf{x}| \left(\nabla p \cdot \dfrac{\mathbf{x}}{|\mathbf{x}|} - ikp \right) = 0,
\end{cases}
\tag{1}
$$

where ∂_n denotes the normal derivative on Γ and i the imaginary unit ($i^2 = -1$). The boundary condition on Γ states that the normal derivative of the total field is zero, which relates to a sound-hard object. The last equation is the Sommerfeld radiation condition that describes the outgoing waves of the unbounded domain.

B. Boundary integral formulation

The Helmholtz system [Eq. (1)] that describes the wave propagation in the exterior domain has a fundamental solution and can therefore be reformulated into a boundary integral equation. Thus instead of solving for a volumetric pressure field, the model is solved for a surface potential. This potential on the scattering surface uniquely determines the scattered field at any point in the exterior domain. The design of the boundary integral formulation typically follows trace theorems. Here a standard approach and nomenclature are used.[33] The formulations are introduced without mathematical rigor; definitions of all operators used in this paper can be found in standard textbooks on BEM.[36,37]

Instead of solving for the scattered field pressure p, the boundary model is written in terms of the unknown surface potential

$$
\varphi = (p + p_{\text{inc}})|_\Gamma,
\tag{2}
$$

representing the restriction of the total pressure field to the surface. The pressure of the wave field scattered into the exterior is given by $p = \mathcal{D}(\varphi)$, where the double-layer potential operator $\mathcal{D}$ is defined as

$$
(\mathcal{D}\varphi)(\mathbf{x}) = \int_\Gamma \partial_{n(\mathbf{y})} G(\mathbf{x}, \mathbf{y}) \varphi(\mathbf{y}) \, d\Gamma(\mathbf{y}), \quad \mathbf{x} \in \Omega_{\text{ext}},
\tag{3}
$$

where the Green's function

$$
G(\mathbf{x}, \mathbf{y}) = \frac{1}{4\pi} \frac{e^{ik|\mathbf{x}-\mathbf{y}|}}{|\mathbf{x} - \mathbf{y}|}, \quad \mathbf{x} \neq \mathbf{y}
\tag{4}
$$

is the fundamental solution of the 3D Helmholtz system [Eq. (1)]. Notice that the double-layer potential operator maps from the potential on the surface to the pressure field in the exterior volume. Now the boundary integral formulation can be obtained with the use of the Dirichlet and Neumann traces of the double-layer potential operator in combination with the boundary conditions. The two different traces result in two independent boundary integral equations, i.e.,

$$
\left(\frac{1}{2} I + M \right) \varphi = -p_{\text{inc}}|_\Gamma,
\tag{5a}
$$

$$
D\varphi = -\partial_n p_{\text{inc}}|_\Gamma,
\tag{5b}
$$

where I denotes the identity operator and M and D the double-layer and hypersingular boundary operators, respectively, given by

$$
(M\varphi)(\mathbf{x}) = -\int_\Gamma \partial_{n(\mathbf{y})} G(\mathbf{x}, \mathbf{y}) \varphi(\mathbf{y}) \, d\Gamma(\mathbf{y}), \quad \mathbf{x} \in \Gamma;
\tag{6a}
$$

$$
(D\varphi)(\mathbf{x}) = -\partial_{n(\mathbf{x})} \int_\Gamma \partial_{n(\mathbf{y})} G(\mathbf{x}, \mathbf{y}) \varphi(\mathbf{y}) \, d\Gamma(\mathbf{y}), \quad \mathbf{x} \in \Gamma.
\tag{6b}
$$

The two boundary integral Eqs. (5a) and (5b) can be used independently to solve the scattering problem. However, both have a null-space consisting of resonance modes that will spoil the computational method with spurious solutions. This can be prevented by considering a linear combination of the equations. That is, for $\eta \in \mathbb{C}$ denoting a coupling parameter with $\Im(\eta) \neq 0$, the Burton-Miller formulation reads

$$
\left(\frac{1}{2} I + M + \eta D \right) \varphi = -p_{\text{inc}}|_\Gamma - \eta \partial_n p_{\text{inc}}|_\Gamma \quad \text{on } \Gamma,
\tag{7}
$$

which has a unique solution.[30] A standard choice of $\eta = -i/k$ has been used throughout.[38]

C. Boundary element discretization

A Galerkin method will be used for the numerical discretization of the boundary integral Eq. (7). The surface Γ is

partitioned into flat triangular elements. The discrete weak formulation is given by

$$\sum_{j=1}^{N}\alpha_j\left\langle\phi_k,\left(\frac{1}{2}I+M+\eta D\right)\phi_j\right\rangle$$
$$=-\langle\phi_k,p_{\text{inc}}|_\Gamma\rangle-\eta\langle\phi_k,\partial_n p_{\text{inc}}|_\Gamma\rangle \quad \text{for } k=1,2,3,\ldots,N,$$

(8)

where $\alpha_j \in \mathbb{C}$ denote the unknown coefficients and ϕ_j the basis functions. The inner product $\langle\cdot,\cdot\rangle$ is given by the standard L_2 inner product. The test and basis functions are given by linear piecewise continuous functions on the triangular surface mesh.

D. Incident HIFU field

The Burton–Miller formulation [Eq. (8)] requires the evaluation of the Cauchy data $(p_{\text{inc}}|_\Gamma, \partial_n p_{\text{inc}}|_\Gamma)$ as excitation for the scattering problem. That is, the pressure field and its normal gradient at the scatterer surface have to be specified. For the accurate simulation of HIFU techniques in medical technology, this requires the modeling of ultrasound transducers. These transducers can be designed such that the transmitted acoustic field is targeted in a prescribed region. In the case of transcostal HIFU treatment, the use of a multi-element array of ultrasound transducers has significant advantages over single-element devices because of their beam-forming capabilities.[39] It is likely that the configuration of the multi-element transducer array has to be optimized for each clinical treatment separately. Here a configuration is being used that has already been optimized for transcostal HIFU application for the same ribcage geometry.[27] More precisely, the transducer array is part of a spherical bowl with a radius of 18 cm. The 256 piston elements are located on a spherical strip with 4 cm aperture diameter and 16 cm array diameter around the negative z axis. All piston elements are assumed to be of the same characteristics. Each piston element is modeled as a disk facing toward the focal point, which is located at the global origin by convention, see Fig. 1.

The pressure field from each piston element is approximated with a numerical quadrature procedure on the disk. Furthermore, for typical configurations, the distance between the transducer array and focal point is large compared to the wavelength. The pressure field from each source can therefore be modeled as the far-field representation of a spherical point source, which is proportional to the fundamental solution [Eq. (4)]. Each source point transmits a field with a prescribed frequency, denoted by f in Hz. The wavenumber is thus given by $k = 2\pi f/c$ where c denotes the speed of sound in the exterior medium, e.g., 1500 m/s in the case of water. Summarizing, the incident HIFU pressure field of a transducer with amplitude A_{td} with a number of N_{piston} pistons of amplitude A_m is modeled as

$$p_{\text{inc}}(\mathbf{x}) = A_{\text{td}}\sum_{m=1}^{N_{\text{piston}}}A_m\sum_{n=1}^{N_{\text{quad},m}}w_{mn}G(\mathbf{x},\mathbf{r}_{mn}),$$
$$\partial_n p_{\text{inc}}(\mathbf{x}) = A_{\text{td}}\sum_{m=1}^{N_{\text{piston}}}A_m\sum_{n=1}^{N_{\text{quad},m}}w_{mn}\nabla G(\mathbf{x},\mathbf{r}_{mn})\cdot\mathbf{n}(\mathbf{x}),$$
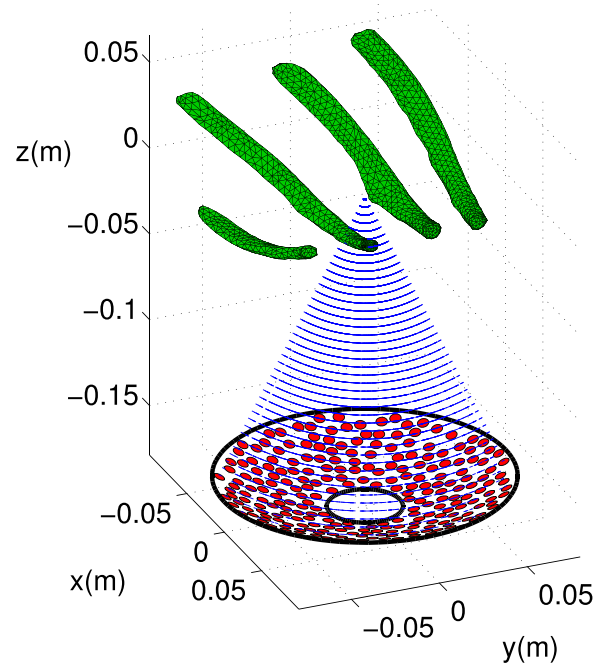
(9)

FIG. 1. (Color online) The model HIFU configuration of a multi-element array of transducers. The transducer array is a spherical strip with 256 piston elements modeled by disks facing to the focal point at the global origin. The ultrasound rays (depicted by the dotted lines) from each piston element focus inside the ribcage. A coarse mesh is visualized on the surface of the ribs.

where $\mathbf{r}_{mn}$ and w_{mn} denote the location and weight of a quadrature point, with $\sum_{n=1}^{N_{\text{quad},m}}w_{mn}=1$ for each piston element m. The quadrature procedure on each piston element is given by a square grid intersected by the disk with equal weights for each point. The amplitudes of the pistons are given by

$$A_m = f\pi a^2 \quad \text{for } m=1,2,\ldots,N_{\text{piston}},$$

(10)

where a denotes the radius of the piston elements, typically 3 mm. The amplitude of the transducer is given by

$$A_{\text{td}} = 2\pi i\rho,$$

(11)

where ρ denotes the mass-density of the exterior medium, e.g., 1000 kg/m^3 for water. For typical simulations, the total number of computational source elements, i.e., the number of pistons times the number of quadrature points, is around 40 000 points.

III. FAST ALGORITHMS FOR HIGH-FREQUENCY SCATTERING

The Burton-Miller formulation [Eq. (7)] is the preferred model equation compared to either the double-layer or the hypersingular integral equation because it has a unique solution and is devoid of spurious modes due to resonances. On the other hand, solving the Burton–Miller formulation with an iterative linear solver such as GMRES requires two matrix-vector multiplications at each iteration and the convergence can be slow. The strongly singular and non-compact hypersingular boundary operator within the

J. Acoust. Soc. Am. **138** (5), November 2015

van 't Wout *et al.*   2729

Burton–Miller formulation results in an unfavorable spectrum because no eigenvalue clustering can be expected for these kinds of operators. The convergence of the linear solver for the standard Burton–Miller formulation is therefore slow, or, even worse, the algorithm might not converge at all within a reasonable amount of iterations. This is especially the case at high frequencies, where the wavelength is small compared to the dimension of the scatterer. High-frequency preconditioners can improve the convergence significantly when they are specifically designed for the particular application.

Next to reducing the number of iterations in the linear solver with preconditioning, the efficiency of the BEM heavily depends on the assembly and storage of the discretization matrices as well. Because the matrices are fully populated, a dense storage requires in excess of hundreds of GB of memory for large-scale scattering problems. For standard computing facilities these simulations are therefore only feasible when the matrices are compressed and stored in a data-sparse format.

A. Operator preconditioning

The Burton–Miller formulation [Eq. (7)] is a linear combination of a double-layer and a hypersingular boundary operator. The hypersingular operator D is an unbounded operator that reduces the regularity of the surface potential. This causes an unfavorable spectrum and therefore slow convergence of the iterative linear solver. The basic idea of preconditioning is to solve a system $V^{-1}D\varphi = V^{-1}f$ instead of $D\varphi = f$ where the operator V has to be designed such that $V^{-1}D$ has a favorable spectrum compared to D and $Vx = b$ relatively easy to solve. In practice, preconditioners are designed according to approximate information on the original operator D. Where the focus of linear preconditioning is on the discrete system, operator preconditioning uses information on the underlying integro-differential equation.[40] These preconditioners typically take the form of boundary operators that approximate the inverse of the continuous model and can achieve beneficial mapping properties between the integro-differential operators. A major advantage of operator preconditioning is its matrix-free design that allows for a seamless combination with acceleration algorithms such as compression techniques. In this paper, the design of the preconditioner will be performed with OSRC techniques.[33] Although not stated explicitly, all other operators are mass-preconditioned,[41] that is, preconditioned with the discretized identity operator.

1. OSRC preconditioner

One of the fundaments of the design of OSRC preconditioners is the use of the Neumann-to-Dirichlet (NtD) map. The corresponding pseudodifferential operator V_{NtD} maps the normal gradient of the pressure field restricted to the surface to the pressure field on the surface, that is,

$$V_{\text{NtD}}(\partial_n p|_\Gamma) = p|_\Gamma, \tag{12}$$

where the pressure field satisfies the exterior Helmholtz model [Eq. (1)]. An important property of the exact NtD map is the relation

$$\frac{1}{2}I + M - V_{\text{NtD}}D = I \quad \text{on } \Gamma. \tag{13}$$

This demonstrates that the exact NtD map is a perfect preconditioner for the Burton–Miller formulation because it yields the identity operator. However, no closed-form expression for the exact NtD map is available, and numerically computing it is as expensive as solving the original Burton–Miller formulation. Instead a high-frequency approximation of the NtD map will be used as preconditioner. This is effective for transcostal HIFU applications because they typically operate in the MHz range, and the wavelength is therefore small compared to ribcages. The design of the high-frequency approximation follows the OSRC method,[31,32] where only the dominant terms of the NtD map in the high-frequency regime are used. The expression of the OSRC preconditioner, denoted by $\tilde{V}$, is given by the pseudodifferential operator

$$\tilde{V} = \frac{1}{ik}\left(1 + \frac{\Delta_\Gamma}{k_\epsilon^2}\right)^{-1/2}, \tag{14}$$

where $k_\epsilon = k + i\epsilon$ for $\epsilon > 0$ denotes a damped wavenumber and Δ_Γ denotes the Laplace–Beltrami operator. The inclusion of damping prevents the occurrence of singularities. A typical choice of damping factor is $\epsilon = 0.4k^{1/3}R^{-2/3}$ where R denotes the radius of the scatterer.[33] The preconditioner $\tilde{V}$ is a regularizing operator and is included in the Burton–Miller formulation as $\eta = -\tilde{V}$. The OSRC-preconditioned Burton–Miller formulation then reads

$$\left(\frac{1}{2}I + M - \tilde{V}D\right)\varphi = -p_{\text{inc}}|_\Gamma + \tilde{V}\partial_n p_{\text{inc}}|_\Gamma \quad \text{on } \Gamma. \tag{15}$$

The preconditioned equation has a unique solution and is a second-kind Fredholm integral equation.[33] Furthermore it satisfies the correct mapping properties for operator preconditioning.[32]

2. Discretization of the preconditioner

The OSRC preconditioner for the Burton–Miller formulation has been designed on a continuous level and resulted in the operator $\tilde{V}$ defined in Eq. (14). The pseudodifferential character of this operator does not allow for a straightforward discretization to obtain a preconditioner in the form of a matrix. Instead a discretized version of the operation $\tilde{V}f$ for a known function f will be developed; this is sufficient for the application of the preconditioner in an iterative linear solver.

To obtain a discrete formulation of the OSRC preconditioner, the square-root operation will be approximated with a Padé series expansion.[42] That is, the equation $u = \tilde{V}f$ is equivalent to

$$ik\left(1+\frac{\Delta_\Gamma}{k_\epsilon^2}\right)u = \left(1+\frac{\Delta_\Gamma}{k_\epsilon^2}\right)^{1/2}f$$

$$\approx C_0 f + \sum_{p=1}^{N_{\text{Padé}}} \frac{\alpha_p}{k_\epsilon^2}\Delta_\Gamma\left(1+\frac{\beta_p}{k_\epsilon^2}\Delta_\Gamma\right)^{-1}f$$

$$(16)$$

for complex-valued coefficients C_0, α_p, β_p. These coefficients are computed with a branch cut with nonzero angle θ because this increases the accuracy of the Padé approximation for the OSRC preconditioner.[43] The preconditioner step within the iterative linear solver is given by $u = \tilde{V}f$ for given vector f and can be expressed in the following sequential way:

$$\left(1+\frac{\beta_p}{k_\epsilon^2}\Delta_\Gamma\right)v_p = f \quad \text{for} \quad p = 1, 2, \dots, N_{\text{Padé}}, \quad (17\text{a})$$

$$g = C_0 f + \frac{1}{k_\epsilon^2}\Delta_\Gamma \sum_{p=1}^{N_{\text{Padé}}} \alpha_p v_p, \quad (17\text{b})$$

$$\left(1+\frac{1}{k_\epsilon^2}\Delta_\Gamma\right)u = \frac{1}{ik}g. \quad (17\text{c})$$

First, N_p Helmholtz equations have to be solved with complex-valued wavenumbers. Then these solutions v_p are combined as a Padé series. Finally, another Helmholtz equation with complex-valued wavenumber has to be solved. This set of $N_{\text{Padé}} + 1$ complex Helmholtz equations can readily be discretized with the same Galerkin method as used for the Burton–Miller formulation. The local character of the Laplace–Beltrami operator yields sparse matrices. A sparse LU factorization will be computed and used to solve the set of Helmholtz systems within each iteration of the linear solver of the preconditioned Burton–Miller formulation. The sparsity of the preconditioner yields very efficient solution algorithms compared to the dense model equation and thus little overhead in both computation time and storage.

B. Compressed storage of matrix

The boundary integral operators of the Burton–Miller formulation are non-local operators and their Galerkin discretization therefore result in fully populated matrices. The storage of these dense matrices requires memory for $\mathcal{O}(N^2)$ floating-point numbers, where N denotes the number of spatial degrees of freedom. For large-scale structures and a high frequency f, the BEM requires $\mathcal{O}(f^2)$ surface elements to represent the wave propagation and thus a storage of $\mathcal{O}(f^4)$, which is prohibitively expensive for most computing platforms. For example, typical HIFU simulations would require in excess of 100 GB RAM storage. The memory footprint of the BEM can be reduced with storage of the matrices in compressed format using data-sparse structures of the discrete system.

The discretization of the boundary integral formulation of the Helmholtz equation results in hierarchical matrices, commonly denoted as $\mathcal{H}$-matrices. The hierarchical structure

of the matrices originates from the Green's function [Eq. (4)], which characterizes the kernels of the double-layer and hypersingular boundary operators [Eq. (6)]. The regularity of the Green's function increases with the distance between the source and observer. Elements of the matrix that correspond to far interactions are therefore more regular than near interactions. This increased regularity of the kernel allows for the accurate approximation with smooth functions. Using these properties for groups of elements result in low-rank approximation of blocks of the matrix, which is the fundament of the $\mathcal{H}$-matrix compression technique.[44,45]

An established version of the $\mathcal{H}$-matrix compression technique will be used, implemented in the library AHMED,[45] which can be used within version 2 of the open-source BEM++ package.[35]

IV. NUMERICAL EXPERIMENTS

The aim of the presented research was to develop an efficient BEM for the scattering analysis of transcostal HIFU modalities in medical treatment of cancer. To this end, the efficiency of the standard Burton–Miller formulation has been improved with OSRC preconditioning and $\mathcal{H}$-matrix compression. In this section, the performance of the algorithm will be assessed with computational experiments. First, the feasibility of the BEM for HIFU modalities will be demonstrated with a simulation of acoustic scattering at 1 MHz on a human ribcage. As benchmark, a comparison with proprietary software for scattering of a ribcage will be performed in Sec. IV B. The improved convergence of the OSRC preconditioner will be demonstrated on a simple sphere model in Sec. IV C. Finally, the effectiveness of the $\mathcal{H}$-matrix compression will be shown in Sec. IV D.

The presented fast BEM formulation has been implemented with the open-source BEM++ library,[35] version 2.0.1. The $\mathcal{H}$-matrix compression of the discrete boundary operators is readily available in BEM++ and is, in this version, based on the AHMED library.[45] The OSRC preconditioner has been implemented within the Python interface of the core C++ library. The GMRES algorithm of the SciPy library has been used as the linear solver for the discrete system, where the termination criterion is a relative error of 10^{-5} in the preconditioned residual with the Euclidean norm. All experiments are performed on a desktop machine with 12 processors and 80 GB RAM.

A. Feasibility of the BEM for HIFU scattering at a ribcage

To demonstrate the feasibility of our fast BEM for transcostal HIFU simulations, let us consider a model of four ribs and a transducer array of 1 MHz. The parameters for the lossless exterior medium are a density of $1000\,\text{kg/m}^3$ and a speed of sound of $1500\,\text{m/s}$, which are characteristic for water. The transducer array consists of 256 piston elements in a configuration that has been optimized for transcostal HIFU simulation, as explained in Sec. II D. The transmitted acoustic field [Eq. (9)] with a wavelength is 1.5 mm focuses inside the ribcage. The scatterer geometry was obtained from medical images of a human ribcage[29] with the length

J. Acoust. Soc. Am. **138** (5), November 2015

van 't Wout *et al.* 2731

of the ribs approximately 12 cm. The boundaries of the sound-hard ribs are partitioned into a surface mesh consisting of 156 575 triangles where the maximum diameter of the triangular elements is 0.5 mm. This results in 78 297 degrees of freedom for continuous piecewise linear test and basis functions. The parameters for the OSRC preconditioner [Eq. (14)] are given by $N_{\text{Padé}} = 2$, $\theta = \pi/3$, and $R = 0.1768$.

The computation of the discrete matrices with the standard $\mathcal{H}$-matrix compression technique took 116 min and the computation of the Cauchy data of the HIFU transducer 9 min. The GMRES solver for the OSRC-preconditioned BEM converged in 94 iterations and 2 min, which is a major improvement over the 4741 iterations and 69 min for the standard Burton–Miller formulation. Figure 2 depicts the pressure field around the ribcage with the HIFU transducer focusing just behind the ribs. The reflections at the ribcage result in a slight distortion of the focused acoustic area.

B. Benchmark comparison

To assess the validity of the fast BEM developed in this paper, a benchmark comparison has been performed with the proprietary program SCATTER provided by PACSYS Limited, which has been used in earlier studies on HIFU scattering as well.[27–29] As test case, let us consider a geometry of three ribs, truncated from the human ribcage model of Fig. 2. The two media are modeled as described before, i.e., sound-hard ribs and water in the exterior. As incident pressure field, let us consider a plane wave $p_{\text{inc}}(\mathbf{r}) = e^{ik\mathbf{r}\cdot\hat{\mathbf{z}}}$ with a frequency of 1 MHz.

The scattering analysis with the BEM presented in his paper has been performed with the BEM++ implementation of the preconditioned Burton–Miller formulation [Eq. (15)] and the standard computational parameters as described before. The Galerkin discretization uses continuous piecewise linear test and basis functions on a triangular mesh with

a maximum meshwidth of 0.375 mm, which results in 57 430 degrees of freedom.

The acoustic simulations performed with SCATTER use a double-layer Helmholtz model [Eq. (5a)] and a complex wavenumber with a small attenuation parameter of $25 \cdot 10^{-15} f^2$. The numerical discretization follows a collocation method with quadratic basis functions on a six-noded triangular surface mesh. The same triangular patches as for BEM++ has been used, resulting in a larger system of 229 692 degrees of freedom for SCATTER. A prescribed number of 40 iterations has been performed for the GMRES linear solver without compression techniques and distributed on a computing cluster of 100 cores.

The computational results depicted in Fig. 3 show a good qualitative agreement of the two algorithms. The BEM++ solver with OSRC preconditioner and $\mathcal{H}$-matrix compression converges within 47 iterations. The computation time for the matrix-fill is 102 min and for the linear solver only 33 s. The prescribed number of 40 GMRES iterations with SCATTER each took approximately 10 min, so in total more than 6 hr.
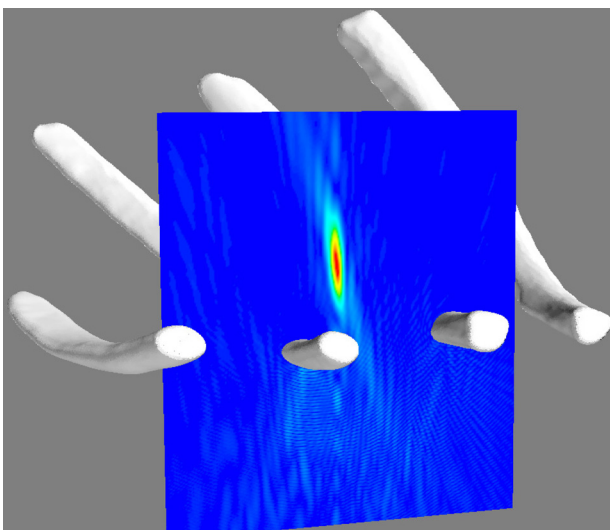

FIG. 2. (Color online) The pressure of the 1 MHz HIFU field focused inside a ribcage, computed with the fast BEM algorithm that uses OSRC preconditioning and $\mathcal{H}$-matrix compression. The transducer array is located below the ribcage.
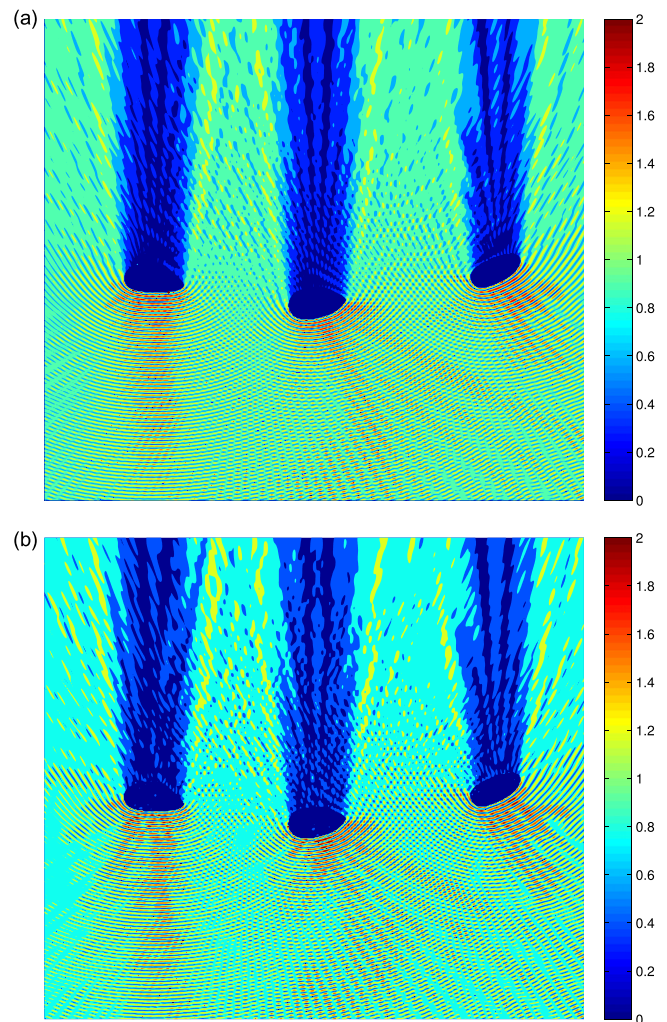

FIG. 3. (Color online) The pressure of a plane wave field propagating in upward direction and scattered at a ribcage visualized on a vertical slice. (a) Preconditioned Galerkin BEM with BEM++. (b) Collocation BEM with SCATTER.

C. Convergence analysis of the preconditioner

The convergence behavior of the OSRC-preconditioned BEM with respect to the frequency has been investigated with a simple model of a unit-sized sphere and an incident plane wave $p_{\text{inc}}(\mathbf{r}) = e^{ik\mathbf{r}\cdot\hat{\mathbf{z}}}$. A convergence analysis has been performed on a set of meshes with maximum element sizes of $\{0.4, 0.2, 0.1, 0.05, 0.025, 0.0125\}$ and corresponding wavelengths $\{8, 4, 2, 1, 0.5, 0.25\}$. This results in a constant oversampling factor of 20, i.e., the wavelength covers at least 20 patches. Standard values for the other computational parameters are used, as described before.

Figure 4 depicts the number of iterations and solution time of the iterative linear solver in the cases of the Burton–Miller formulation and the OSRC-preconditioned formulation with $N_{\text{Padé}} = 2$ and 8. The number of iterations for the preconditioned formulation is almost independent of the frequency, whereas the number of iterations for the standard Burton–Miller formulation rapidly grows with the frequency. This confirms the improved performance of the preconditioner for high-frequency problems. Indeed, each iteration for the preconditioned system is more expensive than the standard Burton–Miller formulation. Still, the wall-clock time of the GMRES solver is smaller for the OSRC-preconditioned formulation. The improvement in computation time is especially significant for high frequencies.
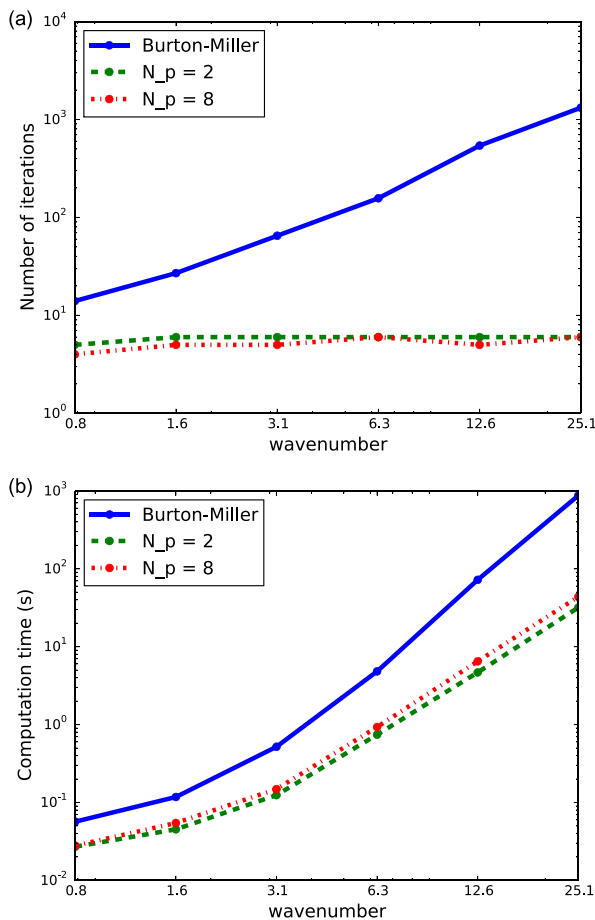

FIG. 4. (Color online) The performance of the Burton–Miller and OSRC-preconditioned formulation on a sphere. The value of N_p denotes the size of the Padé series expansion [Eq. (16)]. (a) The number of GMRES iterations. (b) The simulation time.

TABLE I. The spectral condition numbers of the discretization matrices have been computed for different model formulations of acoustic scattering at a sphere. A small condition number typically results in quick convergence of the iterative linear solver.

Wavenumber	0.79	1.57	3.14	6.28
Burton–Miller	10.81	31.66	90.23	62.10
$N_{\text{Padé}} = 2$	1.63	2.25	2.05	2.04
$N_{\text{Padé}} = 8$	1.15	1.54	1.38	1.38

These convergence characteristics are corroborated by the spectral condition number of the discretization matrix, i.e., the ratio of the largest and smallest magnitude of the eigenvalues. A larger condition number typically results in slower convergence of the iterative linear solver. As listed in Table I, the OSRC preconditioner improves the conditioning of the Burton–Miller formulation. Furthermore, the spectra of the different formulations depicted in Fig. 5 demonstrate that the spectrum of the OSRC-preconditioned equation is more clustered and thus favorable over the Burton–Miller formulation.

One of the most influential parameters for the OSRC preconditioner is the size of the Padé approximation [Eq. (16)]. A larger size of the Padé series improves the accuracy and reduces the number of iterations but increases the computation time. The experiments demonstrate that a relatively small size of the Padé series is already sufficient for the expected performance of the preconditioner.

D. Compression rate of discrete systems

The $\mathcal{H}$-matrix compression technique is based on sparse approximations of low-rank blocks within a matrix. The structure of compressed matrices can be visualized as in Fig. 6. The numbers denote the rank of the approximation of the corresponding block. The compression rate is defined by the ratio of the storage in compressed format compared to the dense matrix. This is related to both the number of blocks and the rank of the approximations.

The $\mathcal{H}$-matrix structure of the discrete double-layer boundary operator for the HIFU simulation described in Sec. IV A is depicted in Fig. 6. The storage of the 78 297 degrees of freedom as a dense matrix would need 93.5 GB RAM. In compressed format it requires 7.4 GB only, which is a compression rate of 7.9%. The algorithm does require some storage overhead during the assembly of the compressed matrix though. For the discrete hypersingular boundary operator, a compression rate of 12.2% to 11.4 GB has been achieved with a structure similar to the double-layer operator. In general, the double-layer operator can be compressed more effectively than the hypersingular operator because the operator is more regular.

1. Influence of frequency and mesh size

The frequency of the wave model has a profound influence on the performance of the compression rate of the $\mathcal{H}$-matrix technique for BEM. A higher frequency typically results in approximations of higher rank and therefore more storage. Simulating a frequency of 0.1 MHz instead of 1 MHz on the same mesh results in a compression rate of

J. Acoust. Soc. Am. **138** (5), November 2015
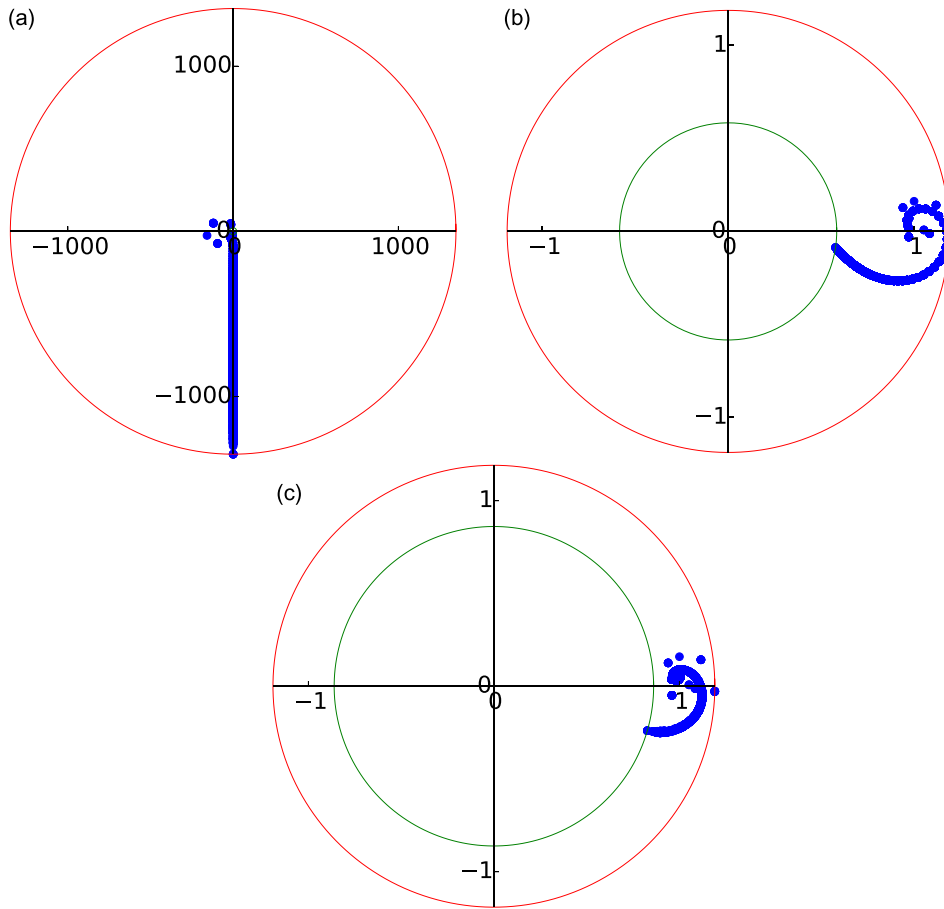
van 't Wout *et al.* 2733

FIG. 5. (Color online) The spectrum of the discrete model equations for the test case on a sphere with a wavenumber of 6.28. The dots depict the eigenvalues of the matrix in the complex plane. The inner and outer circle depict the minimum and maximum value of the spectrum and its bandwidth thus represents the condition number. (a) Burton-Miller. (b) $N_{Padé} = 2$. (c) $N_{Padé} = 8$.
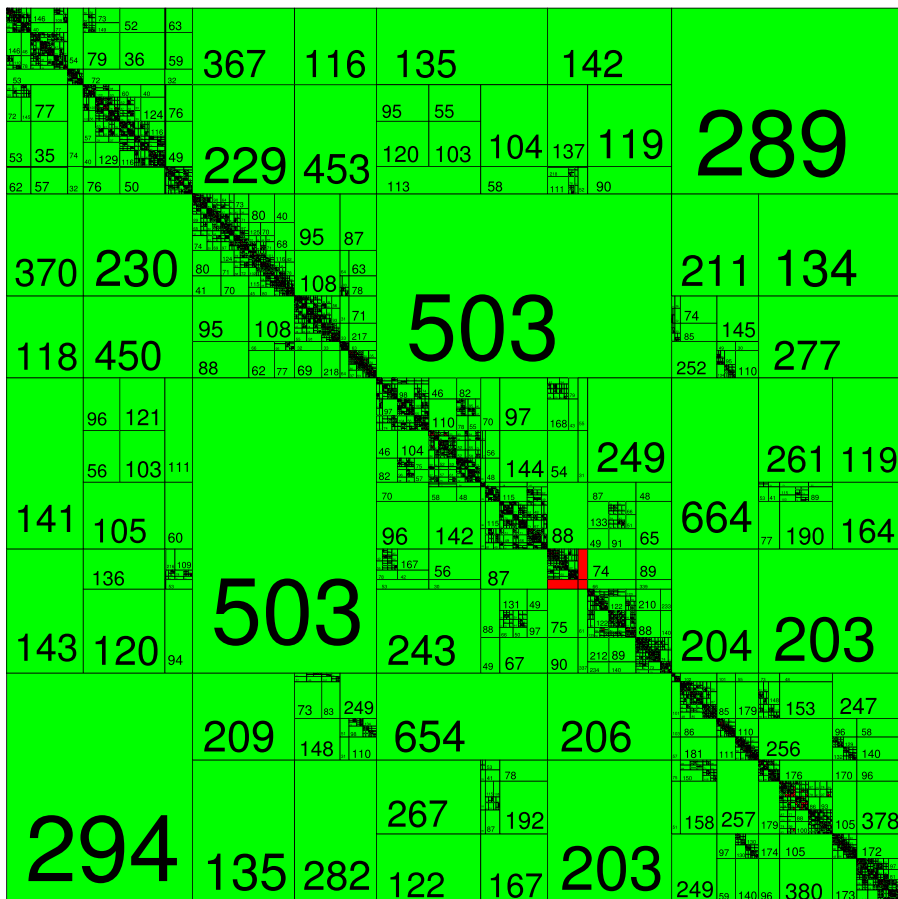

FIG. 6. (Color online) The structure of the data-sparse representation of the discretized double-layer boundary operator with $\mathcal{H}$-matrix compression techniques, as used for the 1 MHz HIFU scattering analysis of a ribcage.
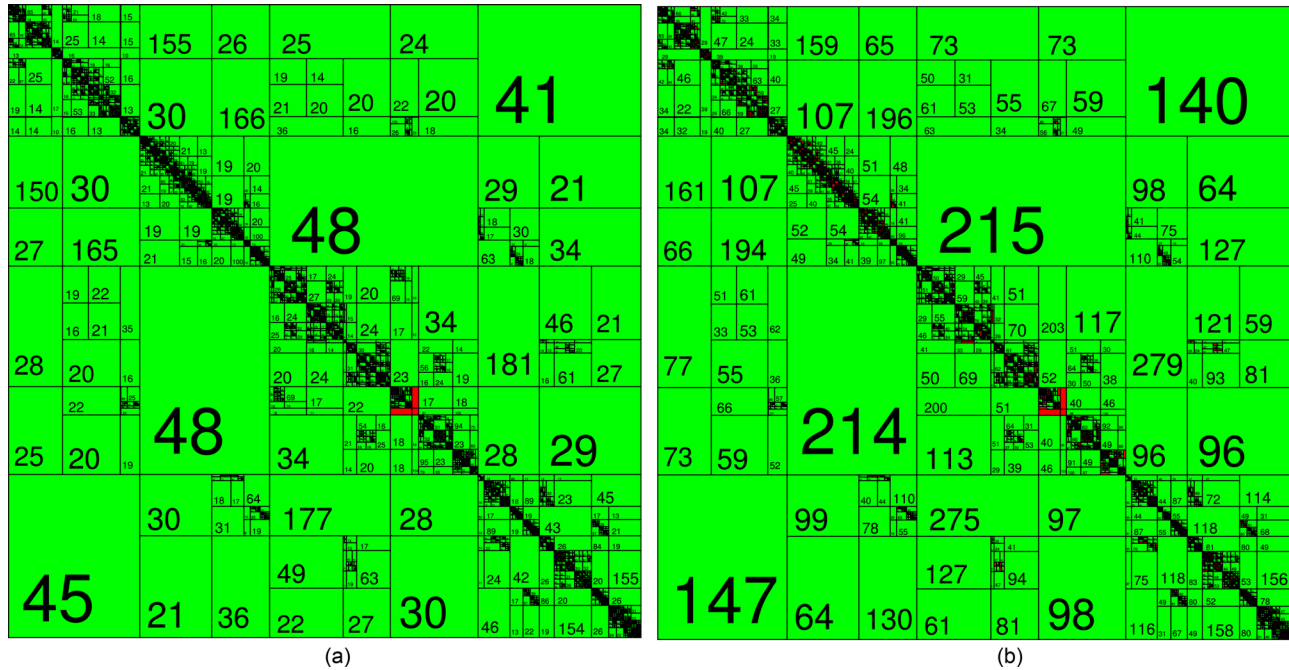
FIG. 7. (Color online) The influence of a lower frequency and a coarser mesh on the structure of the $\mathcal{H}$-matrix compression technique, as compared to Fig. 6. (a) Lower frequency. (b) Coarser mesh.

2.4% to 2.2 GB storage, instead of 7.9%. The structure, depicted in Fig. 7(a), is almost the same for the different frequencies. The rank of the compressed blocks, however, are significantly lower. For instance, the maximum rank drops from 664 to 181 when decreasing the frequency from 1 to 0.1 MHz, respectively.

For BEM, the storage of the dense discrete system scales quadratically with the number of mesh elements. In compressed format, the storage of the matrices may scale less than quadratically when the compression rate increases. To investigate this effect for HIFU simulations, let us

consider the ribcage model with a coarse mesh where the diameter of the elements is two times larger than for the original model. The number of degrees of freedom is 20 600 and the frequency is 0.5 MHz to keep the same oversampling factor. The compression rate of the coarse mesh is only 14.2% compared to the 7.9% for the fine mesh.

2. Influence of geometry

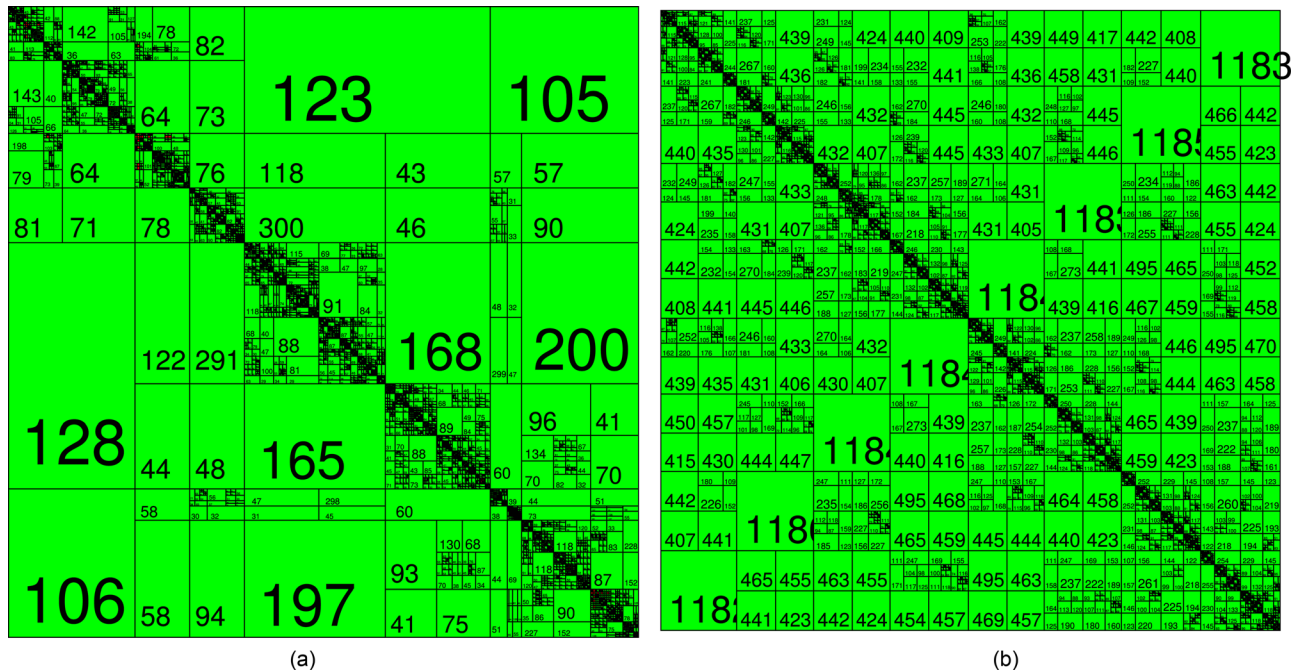The regularity of the Green's function for the 3D Helmholtz equation strongly depends on the distance

FIG. 8. (Color online) The influence of the geometry on the structure of the $\mathcal{H}$-matrix compression technique as compared to Fig. 6. (a) Three rib model. (b) Sphere.

J. Acoust. Soc. Am. **138** (5), November 2015

van 't Wout *et al.*   2735

between spatial elements. The structure of the scatterer can therefore have a large impact on the compression rate of $\mathcal{H}$-matrix techniques for BEM. Specifically, the disjoint ribs in the HIFU model have elements that are far away from each other, relative to the wavelength. The blocks in the discretization matrix corresponding to the influence of the separate ribs can therefore be effectively compressed with a low-rank approximation. This is demonstrated by the $\mathcal{H}$-matrix structure of the four rib model in Fig. 6, where four large blocks can be distinguished. As comparison, let us consider a model of three ribs, used in Sec. IV B as well, and a sphere that consist of 57 430 and 65 187 degrees of freedom, respectively. All simulations are for a frequency of 1 MHz where the sphere is scaled to a diameter of 10 cm for comparable dimensions as the ribcage models. The compression rate for the three-ribs model is 6.3% and for the sphere 22.5%. As depicted in Fig. 8(a), the three ribs result in blocks of one-third the size of the matrix, representing the disjoint regions in the model. The $\mathcal{H}$-matrix compression for the ribcages is significantly more effective than for the sphere. In the case of the sphere, the blocks are relatively small and the rank of the approximations high as depicted in Fig. 8(b). The compression of the discrete system for ribcages with $\mathcal{H}$-matrix techniques is thus more efficient than one might expect from standard test cases such as a sphere.

V. CONCLUSION

The use of HIFU techniques in non-invasive medical therapies can have a large clinical impact on the treatment of a wide range of cancers. The patient-specific planning of transcostal HIFU treatment is likely to rely on numerical simulations to optimize the configuration of the multi-element ultrasound transducer array. A model of perfectly rigid ribs from the human ribcage immersed in an infinite homogenous region of water has been used. Whilst this methodology does not preclude defining locally reacting ribs, the aim of this study was to investigate a rudimentary configuration, which will then be built upon as part of future work. Such models can be efficiently solved using the BEM, which has, however, so far seen limited use due to the high demand for computational resources at high frequencies. An innovative fast BEM has been developed in this paper specifically for application to HIFU simulations. The Galerkin-discretized Burton–Miller formulation is an accurate method for scattering that is devoid of spurious resonances. The innovative use of OSRC-preconditioning techniques significantly improves the convergence for high-frequency scattering at large-scale structures. Furthermore, the $\mathcal{H}$-matrix compression technique effectively reduces the large memory footprint. This novel combination of acceleration algorithms for accurate BEM has been implemented within the open-source library BEM++. Scattering analysis of a human ribcage at 1 MHz confirms the improvement of convergence and the effectiveness of matrix compression with the dedicated fast BEM algorithm. Realistic simulations of transcostal HIFU techniques have been achieved with only 2 hr computation time on a desktop machine. This demonstrates

the applicability of fast BEM simulations to medical treatment with transcostal HIFU modalities.

[1]B. C. Yan and J. A. Hart, "Recent developments in liver pathology," Arch. Pathol. Lab. Med. **133**(7), 1078–1086 (2009).

[2]S. Vaezy, M. Andrew, P. Kaczkowski, and L. Crum, "Image-guided acoustic therapy," Ann. Rev. Biomed. Eng. **3**(1), 375–390 (2001).

[3]N. Y. Piorkowsky, "Europe's hepatitis challenge: Defusing the viral time bomb," J. Hepatol. **51**(6), 1068–1073 (2009).

[4]J. F. P. Bridges, G. Gallego, and B. M. Blauvelt, "Controlling liver cancer internationally: A qualitative study of clinicians' perceptions of current public policy needs," Health Res. Policy Syst. **9**(32), 1–8 (2011).

[5]A. Jemal, F. Bray, M. M. Center, J. Ferlay, E. Ward, and D. Forman, "Global cancer statistics," CA Cancer J. Clinic. **61**(2), 69–90 (2011).

[6]C. H. Cha, M. W. Saif, B. H. Yamane, and S. M. Weber, "Hepatocellular carcinoma: Current management," Curr. Probl. Surg. **47**(1), 10–67 (2010).

[7]D. R. Carpizo and M. D'Angelica, "Liver resection for metastatic colorectal cancer in the presence of extrahepatic disease," Lancet Oncol. **10**(8), 801–809 (2009).

[8]J. S. Tomlinson, W. R. Jarnagin, R. P. DeMatteo, Y. Fong, P. Kornprat, M. Gonen, N. Kemeny, M. F. Brennan, L. H. Blumgart, and M. D'Angelica, "Actual 10-year survival after resection of colorectal liver metastases defines cure," J. Clin. Oncol. **25**(29), 4575–4580 (2007).

[9]G. ter Haar, D. Sinnett, and I. Rivens, "High intensity focused ultrasound–a surgical technique for the treatment of discrete liver tumours," Phys. Med. Biol. **34**(11), 1743–1750 (1989).

[10]G. ter Haar and C. Coussios, "High intensity focused ultrasound: Past, present and future," Int. J. Hypertherm. **23**(2), 85–87 (2007).

[11]L. Crum, M. Bailey, J. H. Hwang, V. Khokhlova, and O. Sapozhnikov, "Therapeutic ultrasound: Recent trends and future perspectives," Phys. Proc. **3**(1), 25–34 (2010). Congress on Ultrasonics, Santiago de Chile, 2009.

[12]T. T. Cheung, S. T. Fan, S. C. Chan, K. S. H. Chok, F. S. K. Chu, C. R. Jenkins, R. C. L. Lo, J. Y. Y. Fung, A. C. Y. Chan, W. W. Sharr, S. H. Y. Tsang, W. C. Dai, R. T. P. Poon, and C. M. Lo, "High-intensity focused ultrasound ablation: An effective bridging therapy for hepatocellular carcinoma patients," World J. Gastroenterol. **19**(20), 3083–3089 (2013).

[13]D. Li, J. Kang, and D. C. Madoff, "Locally ablative therapies for primary and metastatic liver cancer," Expert Rev. Anticancer Ther. **14**(8), 931–945 (2014).

[14]F. Wu, Z.-B. Wang, W.-Z. Chen, H. Zhu, J. Bai, J.-Z. Zou, K.-Q. Li, C.-B. Jin, F.-L. Xie, and H.-B. Su, "Extracorporeal high intensity focused ultrasound ablation in the treatment of patients with large hepatocellular carcinoma," Ann. Surg. Oncol. **11**(12), 1061–1069 (2004).

[15]J.-L. Li, X.-Z. Liu, D. Zhang, and X.-F. Gong, "Influence of ribs on the nonlinear sound field of therapeutic ultrasound," Ultrasound Med. Biol. **33**(9), 1413–1420 (2007).

[16]S. Bobkova, L. Gavrilov, V. Khokhlova, A. Shaw, and J. Hand, "Focusing of high-intensity ultrasound through the rib cage using a therapeutic random phased array," Ultrasound Med. Biol. **36**(6), 888–906 (2010).

[17]J.-F. Aubry, K. B. Pauly, C. Moonen, G. ter Haar, M. Ries, R. Salomir, S. Sokka, K. M. Sekins, Y. Shapira, F. Ye, H. Huff-Simonin, M. Eames, A. Hananel, N. Kassell, A. Napoli, J. H. Hwang, F. Wu, L. Zhang, A. Melzer, Y.-S. Kim, and W. M. Gedroyc, "The road to clinical use of high-intensity focused ultrasound for liver cancer: Technical and clinical consensus," J. Ther. Ultrasound **1**(13), 1–7 (2013).

[18]F. A. Duck, Physical Properties of Tissue: A Comprehensive Reference Book (Academic, London, 1990), p. 346.

[19]B. E. Treeby, J. Jaros, A. P. Rendell, and B. T. Cox, "Modeling nonlinear ultrasound propagation in heterogeneous media with power law absorption using a k-space pseudospectral method," J. Acoust. Soc. Am. **131**(6), 4324–4336 (2012).

[20]Y. Y. Botros, E. S. Ebbini, and J. L. Volakis, "Two-step hybrid virtual array ray (var) technique for focusing through the rib cage," IEEE Trans. Ultrason. Ferroelectr. Freq. Contr. **45**(4), 989–1000 (1998).

[21]P. V. Yuldashev, S. M. Shmeleva, S. A. Ilyin, O. A. Sapozhnikov, L. R. Gavrilov, and V. A. Khokhlova, "The role of acoustic nonlinearity in tissue heating behind a rib cage using a high-intensity focused ultrasound phased array," Phys. Med. Biol. **58**(8), 2537–2559 (2013).

[22]T. 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**(6), 3665–3677 (1999).

2736   J. Acoust. Soc. Am. **138** (5), November 2015

van 't Wout et al.

[23]J.-F. Aubry, M. Pernot, F. Marquet, M. Tanter, and M. Fink, "Transcostal high-intensity-focused ultrasound: Ex vivo adaptive focusing feasibility study," Phys. Med. Biol. **53**(11), 2937–2951 (2008).

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

[25]S. Qiao, G. Shen, J. Bai, and Y. Chen, "Transcostal high-intensity focused ultrasound treatment using phased array with geometric correction," J. Acoust. Soc. Am. **134**(2), 1503–1514 (2013).

[26]J. Jaros, A. P. Rendell, and B. E. Treeby, "Full-wave nonlinear ultrasound simulation on distributed clusters with applications in high-intensity focused ultrasound," Int. J. High Perform. C, in press (2015).

[27]P. Gélat, G. ter Haar, and N. Saffari, "Modelling of the acoustic field of a multi-element HIFU array scattered by human ribs," Phys. Med. Biol. **56**(17), 5553–5581 (2011).

[28]P. Gélat, G. ter Haar, and N. Saffari, "The optimization of acoustic fields for ablative therapies of tumours in the upper abdomen," Phys. Med. Biol. **57**(24), 8471–8497 (2012).

[29]P. Gélat, G. ter Haar, and N. Saffari, "A comparison of methods for focusing the field of a HIFU array transducer through human ribs," Phys. Med. Biol. **59**(12), 3139–3171 (2014).

[30]A. J. Burton and G. F. Miller, "The application of integral equation methods to the numerical solution of some exterior boundary-value problems," Proc. R. Soc. Lond. A Math. Phys. Eng. Sci. **323**(1553), 201–210 (1971).

[31]X. Antoine and M. Darbas, "Alternative integral equations for the iterative solution of acoustic scattering problems," Q. J. Mech. Appl. Math. **58**(1), 107–128 (2005).

[32]X. Antoine and M. Darbas, "Generalized combined field integral equations for the iterative solution of the three-dimensional Helmholtz equation," ESAIM Math. Model. Num. Anal. **41**, 147–167 (2007).

[33]M. Darbas, E. Darrigrand, and Y. Lafranche, "Combining analytic preconditioner and fast multipole method for the 3-D Helmholtz equation," J. Comput. Phys. **236**, 289–316 (2013).

[34]N. A. Gumerov and R. Duraiswami, "A broadband fast multipole accelerated boundary element method for the three dimensional Helmholtz equation," J. Acoust. Soc. Am. **125**(1), 191–205 (2009).

[35]W. Śmigaj, T. Betcke, S. Arridge, J. Phillips, and M. Schweiger, "Solving boundary integral problems with BEM++," ACM Trans. Math. Softw. **41**(2), 6:1–6:40 (2015).

[36]J.-C. Nédélec, *Acoustic and Electromagnetic Equations—Integral Representations for Harmonic Problems* (Springer, New York, 2001), p. 316.

[37]O. Steinbach, *Numerical Approximation Methods for Elliptic Boundary Value Problems—Finite and Boundary Elements* (Springer, New York, 2008), p. 386.

[38]R. Kress, "Minimizing the condition number of boundary integral operators in acoustic and electromagnetic scattering," Q. J. Mech. Appl. Math. **38**(2), 323–341 (1985).

[39]L. Gavrilov and J. Hand, "A theoretical assessment of the relative performance of spherical phased arrays for ultrasound surgery," IEEE Trans. Ultrason. Ferroelectr. Freq. Contr. **47**(1), 125–139 (2000).

[40]R. Hiptmair, "Operator preconditioning," Comput. Math. Appl. **52**(5), 699–706 (2006).

[41]R. Kirby, "From functional analysis to iterative methods," SIAM Rev. **52**(2), 269–293 (2010).

[42]F. A. Milinazzo, C. A. Zala, and G. H. Brooke, "Rational square-root approximations for parabolic equation algorithms," J. Acoust. Soc. Am. **101**(2), 760–766 (1997).

[43]X. Antoine, M. Darbas, and Y. Y. Lu, "An improved surface radiation condition for high-frequency acoustic scattering problems," Comput. Methods Appl. Mech. Eng. **195**(33–36), 4060–4074 (2006).

[44]S. Börm, "Efficient numerical methods for non-local operators: $\mathcal{H}^2$-matrix compression, algorithms and analysis," in *Tracts in Mathematics* (European Mathematical Society, Zürich, 2010), Vol. 14, p. 432.

[45]M. Bebendorf, "Hierarchical matrices: A means to efficiently solve elliptic boundary value problems," in *Lecture Notes in Computational Science and Engineering* (Springer, Berlin, 2008), Vol. 63, p. 296.