Accelerating frequency-domain numerical methods for weakly nonlinear focused ultrasound using nested meshes

Samuel P. Groth, Pierre Gélat, Seyyed R. Haqshenas, Nader Saffari, Elwin van 't Wout, Timo Betcke, and Garth N. Wells

Citation: The Journal of the Acoustical Society of America **150**, 441 (2021); doi: 10.1121/10.0005655
View online: https://doi.org/10.1121/10.0005655
View Table of Contents: https://asa.scitation.org/toc/jas/150/1
Published by the Acoustical Society of America

JASA ARTICLE

Accelerating frequency-domain numerical methods for weakly nonlinear focused ultrasound using nested meshes

Samuel P. Groth,[1,a)] Pierre Gélat,[2] Seyyed R. Haqshenas,[2] Nader Saffari,[2] Elwin van 't Wout,[3] Timo Betcke,[4] and Garth N. Wells[1]

[1]*Department of Engineering, University of Cambridge, Cambridge CB2 1PZ, United Kingdom*

[2]*Department of Mechanical Engineering, University College London, London WC1E 7JE, United Kingdom*

[3]*Institute for Mathematical and Computational Engineering, School of Engineering and Faculty of Mathematics, Pontificia Universidad Católica de Chile, Santiago, Chile*

[4]*Department of Mathematics, University College London, London WC1H 0AY, United Kingdom*

ABSTRACT:

The numerical simulation of weakly nonlinear ultrasound is important in treatment planning for focused ultrasound (FUS) therapies. However, the large domain sizes and generation of higher harmonics at the focus make these problems extremely computationally demanding. Numerical methods typically employ a uniform mesh fine enough to resolve the highest harmonic present in the problem, leading to a very large number of degrees of freedom. This paper proposes a more efficient strategy in which each harmonic is approximated on a separate mesh, the size of which is proportional to the wavelength of the harmonic. The increase in resolution required to resolve a smaller wavelength is balanced by a reduction in the domain size. This nested meshing is feasible owing to the increasingly localised nature of higher harmonics near the focus. Numerical experiments are performed for FUS transducers in homogeneous media to determine the size of the meshes required to accurately represent the harmonics. In particular, a fast *volume potential* approach is proposed and employed to perform convergence experiments as the computation domain size is modified. This approach allows each harmonic to be computed via the evaluation of an integral over the domain. Discretising this integral using the midpoint rule allows the computations to be performed rapidly with the FFT. It is shown that at least an order of magnitude reduction in memory consumption and computation time can be achieved with nested meshing. Finally, it is demonstrated how to generalise this approach to inhomogeneous propagation domains. © *2021 Acoustical Society of America*. https://doi.org/10.1121/10.0005655

(Received 9 November 2020; revised 27 June 2021; accepted 1 July 2021; published online 20 July 2021)

[Editor: Nail A. Gumerov]

Pages: 441–453

I. INTRODUCTION

Focused ultrasound (FUS) is a non-invasive tumor ablation therapy in which acoustic waves are focused at a target location, thereby elevating the temperature sufficiently to destroy the tumor tissue. Tissue death can be caused directly via thermal ablation or via other mechanisms, such as cavitation (Izadifar *et al.*, 2017; Vlaisavljevich *et al.*, 2013), shock-scattering histotripsy (Maxwell *et al.*, 2011), and boiling histotripsy (Khokhlova *et al.*, 2011). In the thermal ablation setting, which is the focus of this article, the peak acoustic pressure is often between 1 and 10 MPa, at which linear acoustic theory is no longer accurate. Indeed, it has been observed that the contributions from nonlinear effects can increase the temperature at the focal point by an additional 20%, thus substantially reducing the time required for ablation when compared to simulations using linear theory (Solovchuk *et al.*, 2014). To simulate the nonlinear acoustic field, a popular model is the Westervelt equation, which incorporates a quadratic pressure term. Numerically solving the Westervelt equation in a FUS setting is computationally challenging, since the domains are large compared to the smallest wavelength present.

To put into context the scale of the computational problems presented by FUS, consider a bowl-shaped FUS transducer with radius 7.5 cm and geometric focal length 15 cm operating at 1.1 MHz. Transducers of this size are suitable for deep-seated tumors, located in the liver, for example. A simulation domain spanning the face of the bowl and extending to the focal region is approximately 100 wavelengths across in each dimension. If there are five harmonics present in the nonlinear field, then this is 500 wavelengths at the fifth harmonic. To obtain a reasonable accuracy level, we can assume that at least six degrees of freedom (DOF) per wavelength are required (Marburg, 2008). Then to simulate this problem in three dimensions, we would require 27×10^9 DOF. Handling such a large system presents an enormous computational load in terms of memory and time.

There has been a great deal of research effort aimed at reducing the computational cost of nonlinear FUS simulations. Much of this work employs at least one simplifying assumption, such as axisymmetry, one-way wave propagation, or the parabolic approximation [see Gu and Jing (2015) for a review]. Popular methods for solving the original

a)Electronic mail: samuelpgroth@gmail.com

Westervelt equation (or equivalent) include the finite-difference time-domain method (FDTD) (Solovchuk *et al.*, 2013) and the *k*-space pseudospectral method (KSPS) (Treeby *et al.*, 2012). Although powerful techniques, these full-wave solvers can still yield simulation times in excess of a day for realistic problems when run on a large cluster (Jaros *et al.*, 2016). For efficiently simulating highly focused problems with strong nonlinearities, shock-capturing numerical schemes have been developed and applied for the parabolic approximation of the Westervelt equation, known as the KZK equation (Bessonova *et al.*, 2009), wide-angle parabolic approximation (Yuldashev *et al.*, 2018), and the one-way propagation version of the Westervelt equation (Yuldashev and Khokhlova, 2011).

In this paper, we set out to improve upon the meshing strategies employed in popular full-wave solvers for the Westervelt equation. These solvers typically use uniform meshes resolved according to the wavelength of the highest harmonic. Such a meshing approach is inefficient, since the higher harmonics are localised around the focus; therefore, the extremely fine mesh far away from the focus is likely overkill. It seems intuitive that a mesh that is gradually refined as the focus is approached would be sensible, but at what rate? And what saving can one expect to achieve? Here, we address these questions via numerical experimentation on realistic FUS transducers taken from Sonic Concepts (2020).

We note similar strategies have previously been employed to reduce the computational load of nonlinear acoustics simulations (Karzova *et al.*, 2017; Yuldashev and Khokhlova, 2011). In these prior works, it was also observed that far from the transducer focus, higher harmonics provide negligible contributions to the acoustic field, and therefore, to save on storage requirements, higher harmonics are stored only on smaller domains, centred at the focus. However, the grid step used in these works is not altered for the different harmonic computations. By increasing the grid step for lower harmonics, further computational savings can be achieved, as we demonstrate in this article. Furthermore, we detail an in-depth study of how to choose appropriate computational domains for each harmonic, and we provide a useful rule of thumb relating the domain dimensions to each harmonic's wavelength.

To perform these experiments efficiently, we consider a simplified setting in which the propagation medium is homogeneous, and the signal is assumed to have harmonic time dependence. Furthermore, we consider settings in which the peak amplitude is no more than 15 MPa and thus the field is only weakly nonlinear. Under the time-harmonic and weakly nonlinear assumptions, the full-wave Westervelt equation reduces to a series of inhomogeneous Helmholtz equations, one for each harmonic present in the field [as in, e.g., Du and Jensen (2013)]. We note that the weak nonlinearity permits us to reasonably neglect the transfer of energy from higher to lower harmonics. This is, however, not valid for extremely high amplitude fields, as encountered in histotripsy applications.

Since we consider a homogeneous medium, each Helmholtz equation is exactly solved by the evaluation of the appropriate *volume potential* integral (Costabel, 2015). Therefore, for each harmonic in the field, we look at the convergence of the volume potential as the domain of integration is increased. To efficiently evaluate the volume potential, we discretise the domain over a uniform voxel mesh, which allows the summation (i.e., the discrete form of the integration) to be performed rapidly using the fast-Fourier transform (FFT). This is an extremely efficient technique for computing high-order harmonics in a homogeneous medium, and its application in this area is, to the best of the authors' knowledge, novel.

For numerical investigations with different transducer configurations and within two different media (water and liver), we deduce that (for the configurations considered) to perform accurate computations of the second harmonic, the computation domain must extend all the way back to the transducer; however, it may be contracted slightly in the transverse direction. For the third and higher harmonics, the computation domain may be contracted in both the axial and transverse directions, thus leading to substantial computational savings. In fact, we demonstrate that an accurate approximation (less than 1% relative error) may be obtained while contracting the width and length of the domain in proportion to the wavelength of the harmonic. Thus, the number of cells in every mesh is roughly equal. This amounts to a reduction in the number of DOF by a factor of approximately $(n/2)^3$, where n is the number of harmonics being computed.

In practice, this leads to a series of nested meshes, each at the resolution required for the appropriate harmonic and all with the same number of voxels. To perform computations for higher harmonics, solutions from lower harmonics are interpolated onto the finer meshes. We outline an algorithm for this more efficient computation of all the harmonics in Sec. VI and examine its performance.

The layout of the paper is as follows. Section II outlines the mathematical model we employ for our FUS setup. In particular, we consider the Westervelt equation and review how it reduces to a series of Helmholtz equations under the time-harmonic assumption. We further simplify the equations via an assumption of weak nonlinearity. In Sec. III, we describe how, in the homogeneous domain case, the Helmholtz equations are solved exactly via volume potentials. It is then described how these volume potentials are efficiently approximated using a voxelised discretisation approach. This leads to the discrete versions of the volume potentials having block-Toeplitz form; thus, the potentials may each be evaluated using a single fast matrix-vector product with the FFT. Section III B discusses our model for the time-harmonic bowl-shaped transducer. Section IV presents a validation of our approach by comparing to simulations performed with the HITU Simulator MATLAB toolbox (Soneson, 2017). In Sec. V, we perform convergence tests for each of the harmonics for a range of problem setups. We present our findings on the rates of convergence of the

https://doi.org/10.1121/10.0005655

approximation as the domain is increased and present a rule of thumb for designing the meshes to ensure each harmonic is accurately approximated. In Sec. VI, we present how the hierarchy of meshes is used in practice with our volume potential approach, including the interpolation of approximations between the meshes. Here, we present some performance details for this three-dimension full-wave approach on a single workstation. Finally, in Sec. VIII, we present our conclusions and discuss the relevance of this work to other numerical techniques for FUS.

II. NONLINEAR ACOUSTICS IN THE FREQUENCY DOMAIN

Acoustic fields produced by FUS transducers are commonly modelled by the Westervelt equation (Hamilton and Blackstock, 1998),

$$\nabla^2 p - \frac{1}{c_0^2}\frac{\partial^2 p}{\partial t^2} - \frac{2\alpha_0}{c_0^{1-\eta}}\frac{\partial}{\partial t}(-\nabla^2)^{\eta/2} p = -\frac{\beta}{\rho_0 c_0^4}\frac{\partial^2 p^2}{\partial t^2}, \tag{1}$$

where p is the acoustic pressure, c_0 is the speed of sound, ρ_0 is the medium density, β is the nonlinearity parameter, and α_0 and η are medium specific attenuation parameters. The fractional Laplacian appearing in (1) was first introduced in Chen and Holm (2004) to incorporate frequency-dependent power law attenuation. More specifically, in the frequency domain, it generates a complex wavenumber of the form

$$k = \frac{\omega}{c_0} + \mathrm{i}\alpha; \quad \alpha = \alpha_0 |\omega|^\eta, \tag{2}$$

for $\alpha_0|\omega|^{\eta-1}c_0 < 0.1$ (Szabo, 1994), where ω is the angular frequency of the transducer. The power law exponent is typically in the range $1 \le \eta \le 2$. We note that this power law attenuation can also be incorporated via a temporal convolution, as was originally proposed in Szabo (1994); we refer the reader to Treeby and Cox (2010) for a review of power law attenuation techniques. We note that the relation (2) does not incorporate the effect of dispersion on the real part of the wavenumber. However, for the frequencies and materials considered in this article, this effect is negligibly small. For highly nonlinear and higher frequency problems, the relation may be modified to more accurately incorporate dispersion, as discussed in Treeby and Cox (2010) and Waters *et al.* (2005).

In this article, we assume that the operation time of the transducer is long when compared to the period of the signal. Therefore, the total acoustic pressure can be written as the following sum over harmonics [as in Du and Jensen (2013) and Soneson (2017)]:

$$p(\mathbf{x}, t) = \mathrm{Re}\left\{\sum_{n=1}^{\infty} p_n(\mathbf{x}) e^{-\mathrm{i}n\omega t}\right\}. \tag{3}$$

Many numerical schemes in the literature consider the time-harmonic form (3) (e.g., Campos-Pozuelo *et al.*, 1999; Du and Jensen, 2013; Soneson, 2017; van 't Wout *et al.*, 2015), likely owing to the distinct advantages of a frequency-domain approach:

- The challenging task of choosing/developing an efficient time stepping scheme can be avoided;
- Arbitrary frequency-dependent attenuation power laws can be easily incorporated (whereas in the time domain, one has to contend with the fractional Laplacian);
- Methods such as the boundary element method and volume integral equation (VIE) method can be applied directly;
- The computation regions for higher harmonics can be reduced, thereby making simulations more efficient.

It is this final point that we study in this article.

The form (3) is not well suited for substitution into (1), since the right-hand side of (1) requires the computation of a product. For such a product, expressions in which real or imaginary parts are required to be taken lead to cumbersome algebra. Therefore, it is more straightforward to use the following expression, which is equivalent to (3):

$$p(\mathbf{x}, t) = \frac{1}{2}\sum_{n=1}^{\infty}\left(p_n(\mathbf{x}) e^{-\mathrm{i}n\omega t} + p_n^*(\mathbf{x}) e^{\mathrm{i}n\omega t}\right), \tag{4}$$

where * denotes complex conjugation. Substituting (4) into (1) and matching coefficients of $e^{-\mathrm{i}n\omega t}$ for $n \ge 1$ yields

$$\nabla^2 p_n + k_n^2 p_n = \frac{\beta\omega^2}{2\rho_0 c_0^4} n^2 \sum_{m=1}^{\infty} p_m (p_{n-m} + 2p_{m-n}^*), \tag{5}$$

for $n = 1, 2, \ldots$, where $p_n = 0$ for $n \le 0$, and the complex wavenumbers are defined as

$$k_n = \frac{n\omega}{c_0} + \mathrm{i}\alpha(n\omega). \tag{6}$$

Equation (5) can be further simplified by neglecting small terms on the right-hand side, which is appropriate in the weakly nonlinear case, as was considered in, e.g., Du and Jensen (2013). Specifically, we neglect all terms $p_i p_j$ and $p_i p_j^*$ for which $i + j > n$, for $n = 1, 2, \ldots$, thus giving

$$\nabla^2 p_n + k_n^2 p_n = \frac{\beta\omega^2}{2\rho_0 c_0^4} n^2 \sum_{m=1}^{n-1} p_m p_{n-m}. \tag{7}$$

This is a cascade of inhomogeneous Helmholtz equations in which each right-hand side is a combination of products of lower harmonics. Therefore, we can solve the equations sequentially, starting from $n = 1$.

We note that the assumption of weak nonlinearity is valid for a range of FUS applications for thermal ablation, for example, at a preclinical stage, when excitation protocols and devices require characterisation (Kothapalli *et al.*, 2018; Ries *et al.*, 2010). However, for extremely high focal pressures such as those encountered in lithotripsy and histotripsy, where pressures much higher than 30 MPa can be

seen (Izadifar *et al.*, 2017), the terms neglected above become significant. Therefore, a modified approach would be required; however, we do not consider this case here.

III. VOLUME POTENTIALS

Equations (7) are each of the general form

$$\nabla^2 u + k^2 u = f(\mathbf{x}), \quad \mathbf{x} \in \mathbb{R}^3. \tag{8}$$

Via Green's theorem (see, e.g., Colton and Kress, 2013; Costabel, 2015), it can be seen that the following volume potential satisfies (8):

$$u(\mathbf{x}) = -\int_{\mathbb{R}^3} G_k(\mathbf{x},\mathbf{y}) f(\mathbf{y}) \mathrm{d}\mathbf{y}, \quad \mathbf{x} \in \mathbb{R}^3, \tag{9}$$

where G_k is the fundamental solution to the Helmholtz equation, also known as Green's function

$$G_k(\mathbf{x},\mathbf{y}) = \frac{e^{\mathrm{i}k|\mathbf{x}-\mathbf{y}|}}{4\pi|\mathbf{x}-\mathbf{y}|}, \quad \mathbf{x} \neq \mathbf{y}. \tag{10}$$

This function is singular when $\mathbf{x} = \mathbf{y}$; however, integrals of the function across this singularity may be evaluated using principal value techniques or appropriate coordinate transformations, as we shall see in Sec. III A. We note that the integral in (9) is over an infinite domain; however, in practice, we replace this with a finite domain of integration, D. Thus, we have that

$$u(\mathbf{x}) = -\int_D G_k(\mathbf{x},\mathbf{y}) f(\mathbf{y}) \mathrm{d}\mathbf{y} + \varepsilon(D), \quad \mathbf{x} \in D \subset \mathbb{R}^3, \tag{11}$$

where $\varepsilon(D)$ is an error incurred by the introduction of a finite integration domain. Since the FUS field is highly focused, a sensibly chosen finite D will yield a negligibly small error. It is the purpose of this article to investigate how small D can be made whilst still yielding accurate approximations to u.

For clarity, we write out this integral representation (11) for each of the first five harmonics,

$$p_2(\mathbf{x}) = -\frac{2\beta\omega^2}{\rho_0 c_0^4} \int_{D_2} G_{k_2}(\mathbf{x},\mathbf{y}) p_1^2(\mathbf{y}) \mathrm{d}\mathbf{y}, \tag{12}$$

$$p_3(\mathbf{x}) = -\frac{9\beta\omega^2}{\rho_0 c_0^4} \int_{D_3} G_{k_3}(\mathbf{x},\mathbf{y}) p_1(\mathbf{y}) p_2(\mathbf{y}) \mathrm{d}\mathbf{y}, \tag{13}$$

$$p_4(\mathbf{x}) = -\frac{8\beta\omega^2}{\rho_0 c_0^4} \int_{D_4} G_{k_4}(\mathbf{x},\mathbf{y}) (p_2^2(\mathbf{y}) + 2p_1(\mathbf{y}) p_3(\mathbf{y})) \mathrm{d}\mathbf{y}, \tag{14}$$

$$p_5(\mathbf{x}) = -\frac{25\beta\omega^2}{\rho_0 c_0^4} \int_{D_5} G_{k_5}(\mathbf{x},\mathbf{y}) (p_1(\mathbf{y}) p_4(\mathbf{y}) + p_2(\mathbf{y}) p_3(\mathbf{y})) \mathrm{d}\mathbf{y}, \tag{15}$$

where we have introduced the domains D_i, $i = 2, 3, \ldots$, appropriate sizes of which are to be determined. In a homogeneous medium, the first harmonic p_1 is merely the incident field generated by the transducer, which we can compute anywhere in $\mathbb{R}^3$, as is discussed in Sec. III B. Note finally that the Green's functions in each of (12)–(15) possess the appropriate wavenumber for that harmonic, k_n, for $n = 2, 3, 4, 5$.

A. Efficient computation of the volume potential

FUS problems are renowned for their challenging high-frequency nature, with domain sizes up to hundreds of wavelengths in each of the three dimensions. Therefore, the (singular) integrals in (12)–(15) are potentially enormously expensive to approximate.

We choose the integration domains D_i to be cuboidal in shape and to be discretised into uniform voxel grids $\mathcal{V}(D_i)$ so that the discrete form of the integral operator becomes a Toeplitz matrix. A Toeplitz matrix of dimension N has the property that it may be embedded in a circulant matrix of dimension $2N$, with which a matrix-vector product can be computed with $\mathcal{O}(2N \log 2N)$ complexity using the FFT; see, e.g., Groth *et al.* (2020) for more details.

For the voxels in which the Green's function's singularity is located, we approximate the integral by the integral over a sphere of radius a, chosen such that the sphere's volume is equal to that of the voxel. This integral can then be transformed into spherical coordinates, which is convenient because the Jacobian determinant of the transformation cancels the singularity in the Green's function. In the non-singular voxels, we approximate the volume potential integral by the midpoint rule. This gives us the simple quadrature rule

$$\int_{V_j} G(\mathbf{x}_i,\mathbf{y}) f(\mathbf{y}) \mathrm{d}\mathbf{y} = \begin{cases} \frac{1}{k^2}\left\{e^{\mathrm{i}ka}(1 - \mathrm{i}ka) - 1\right\} f(x_j), & \mathbf{x}_i \in V_j, \\ (\delta x)^3 G(\mathbf{x}_i,\mathbf{x}_j) f(\mathbf{x}_j), & \mathbf{x}_i \notin V_j, \end{cases} \tag{16}$$

for $i, j = 1, \ldots, N$, where N is the number of voxels in the grid and δx is the side length of each voxel V_j. This is reminiscent of the "discrete dipole approximation" often used in electromagnetic scattering calculations (Draine and Flatau, 1994). One can opt for a more sophisticated quadrature rule here; however, this simple approach suffices for our purposes.

B. FUS incident field

In a homogeneous medium, there is no scattering, and therefore the first harmonic, p_1, can be obtained directly from an appropriate model of the time-harmonic field generated by the FUS source. In this paper, we consider bowl-shaped ultrasound transducers, which are designed to focus acoustic energy at a prescribed location, typically the centre

https://doi.org/10.1121/10.0005655

of curvature of the bowl. More specifically, we consider for simplicity a single-element bowl-shaped transducer and discretise the surface using a Rayleigh integral type approach described below. We note that we are not exploiting symmetries in our approach, and therefore more sophisticated multi-element transducer arrays, such as those in Gavrilov and Hand (2000), Gélat *et al.* (2011), and Kreider *et al.* (2013), can be incorporated in a straightforward manner.

To discretise the surface of the bowl transducer, we use evenly spaced points following (Deserno, 2004). At each of the evenly spaced points, we place a monopole source, the expression for which is given by (10). Summing over the monopole sources gives the (unnormalised) first harmonic at any $\mathbf{x} \in \mathbb{R}^3$ as

$$\tilde{p}(\mathbf{x}) = \frac{A}{n_\mathrm{p}} \sum_{i=1}^{n_\mathrm{p}} G_k(\mathbf{x}, \mathbf{r}_i), \tag{17}$$

where n_p is the number of points, $\mathbf{r}_i$ are their locations on the bowl, and A is the total surface area of the bowl.

We note that in (17), no amplitude has been specified for the monopole sources. We instead choose to normalise the field to produce a prescribed total radiated power, Π, which is obtained by integrating the intensity over a sphere surrounding the source (Kinsler *et al.*, 1999). For a bowl-shaped transducer, the field is directed, so rather than integrating over a sphere, it suffices to integrate over a disk covering the open end of the bowl. This can be written as

$$\Pi(p) = \frac{1}{2\rho_0 c_0} \int_0^{2\pi} \int_0^R p^2(r, \theta) \mathrm{d}r\mathrm{d}\theta, \tag{18}$$

where R is the outer radius of the bowl. Thus, the normalised first harmonic to yield a prescribed radiated power Π_0 is given as

$$p_1(\mathbf{x}) = \sqrt{\frac{\Pi_0}{\Pi(\tilde{p})}}\tilde{p}. \tag{19}$$

In our experiments, we take $n_\mathrm{p} = 4096$, which equates to approximately ten monopole sources per fundamental wavelength. Such a large value was required to avoid undesired interference patterns between the bowl and focus.

The two bowl transducer geometries considered in this article are taken from the Sonic Concepts website (Sonic Concepts, 2020) and are detailed in Table I. Two propagation media are considered throughout: water and liver. The acoustic parameters for these are detailed in Table II.

TABLE I. Operating frequencies and geometrical parameters of bowl transducers considered, taken from Sonic Concepts (2020). The geometrical parameters are the geometric focal length l and the outer radius R. These can also be seen depicted in Fig. 1.

	f_0 (MHz)	l (mm)	R (mm)
H101	1.1	63.2	32
H131	1.1	35	16.5

TABLE II. Relevant medium parameters for water, liver, and kidney at 1 MHz (Azhari, 2010; Duck, 2013). The parameters α_0 and η pertain to the absorption power law in Eq. (2).

	ρ_0 (kg/m^3)	c (m/s)	β	α_0	η (dB/m)
Water	1000	1480	3.5	0.2	2
Liver	1060	1590	4.4	90.0	1.1
Kidney	1050	1570	4.7	10	1

C. How many points per wavelength?

Before deploying our scheme to compute high-order harmonics, it is necessary to understand the convergence rate as the mesh is refined, thereby enabling an adequate resolution to be chosen for later investigations. In this article, we focus on achieving approximations with relative L^2 errors close to 1%.

To determine the convergence rate, we consider computing the second harmonic generated by the H131 transducer in liver, using Eq. (12). The configuration of the transducer is shown in Fig. 1; the bowl transducer is located at the origin and is directed along the x axis. For the integration domain D_2 in (12), we take the white box in Fig. 1. This box is illustrated with detailed dimension definitions in Fig. 2. The domain D_2 is defined as

$$D_2 = [l - L, l + d] \times [-R, R] \times [-R, R], \tag{20}$$

where d is the distance of interest beyond the focus, l is the geometric focal length, and R is the outer radius (see Fig. 2). The length L is defined as $L = \sqrt{l^2 - R^2} - \epsilon$ with ϵ chosen as a small displacement from the bowl to avoid the possibility of $\mathbf{x} = \mathbf{r}_i$ in (17), for which the monopole sources are undefined; we take $\epsilon = 0.1$ mm. In this article, we suppose that the region near the focus is of primary interest; therefore, from a computational point of view, we desire to shrink the computation domain to be much shorter and narrower than L and w, i.e., more localised around the focus, to reduce

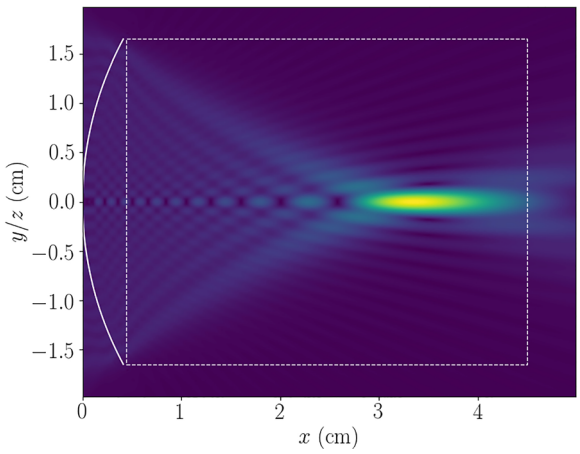

FIG. 1. (Color online) Magnitude of the first harmonic generated by the H131 transducer in liver. The bowl transducer is represented by the arc at the far left. The dashed white line outlines the computation domain D_2 used to compute the second harmonic shown in Fig. 4.

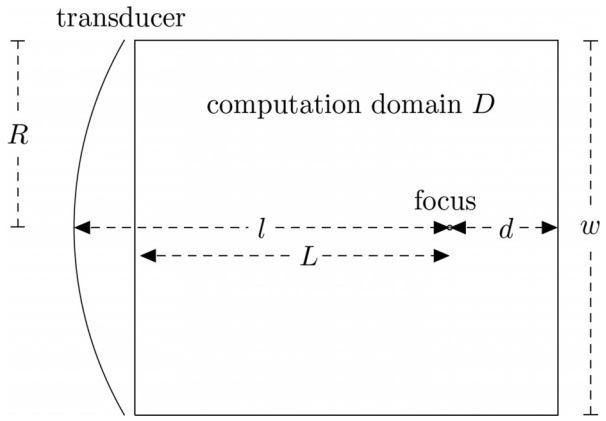

FIG. 2. Description of a generic computation domain D with length $L+d$ and width (and depth) w.

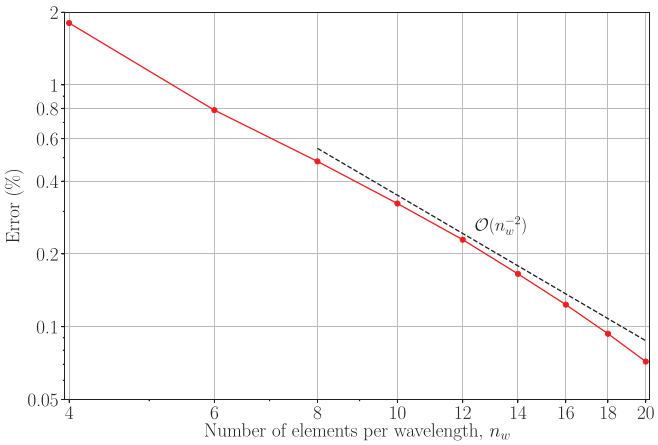

FIG. 3. (Color online) The convergence of the quadrature rule for the computation of the second harmonic via (12). The convergence rate is quadratic in n_w, and an error of smaller than 1% is achieved with $n_w > 5$.

computational load. The extent to which this shrinking can be done without losing accuracy is the focus of Secs. V and VI. The values of L and w defined above represent our default computation domain. The value of the distance d depends on the user's interest in the region beyond the focus. In the absence of scatterers beyond the focus, the field in this region after the focus will not affect the intensity at the focus. In the present example, the total radiated power of the transducer is set as $\Pi_0 = 100\,\mathrm{W}$, and the operating frequency is $f_0 = 1.1\,\mathrm{MHz}$.

The volume potential in (12) is computed using the method outlined in Sec. III A with increasingly refined voxel meshes. Specifically, we build meshes with voxel dimension $\delta x = \lambda/(2n_w)$, where $\lambda = c/f_0$ is the fundamental wavelength and n_w is the "number of voxels per wavelength." The factor 2 in the denominator is to account for the fact that we are computing a field with wavelength $\lambda/2$, i.e., the second harmonic. The values of n_w considered are from 4 to 20, and a reference solution, denoted p_2^r, is computed using $n_w = 35$. For each value of n_w, the relative L^2-error of the field p_2 along the x axis is computed; this is defined as

$$\text{Error} = \frac{||p_2 - p_2^r||}{||p_2^r||} \times 100\%, \tag{21}$$

where $||\cdot||$ denotes the L^2-norm along the x axis, i.e.,

$$||p_2|| := \left(\int_{l-L}^{l+d} |p_2(x,0,0)|^2 \mathrm{d}x \right)^{1/2}. \tag{22}$$

We approximate the integrals in (21) using the midpoint rule with the mesh nodes of the reference solution, p_2^r, being used as the quadrature nodes.

The convergence results are shown in Fig. 3. As is to be expected from the midpoint rule, quadratic convergence is obtained. From the graph, we can read off that an error smaller than 1% is achieved with $n_w > 5$. Therefore, we choose to take $n_w = 6$ for all harmonics in the experiments in the remainder of the article. The approximation to the second harmonic with $n_w = 6$ is shown in Fig. 4 and can be seen to be indistinguishable from the reference solution.

IV. VALIDATION OF NUMERICAL SCHEME

To validate our approach for the computation of higher harmonics via the evaluation of the volume potentials in (12)–(15), we present a qualitative comparison to approximations obtained using HITU Simulator (Soneson, 2017), which we have chosen because of its computational efficiency for the axisymmetric problems considered here. HITU simulator is an open-source MATLAB implementation of the high-order parabolic approximation to the axisymmetric Westervelt equation, i.e., the wide-angle Khokhlov–Zabolotkaya–Kuznetsov (WAKZK) equation. The method is detailed by the author of HITU Simulator (Soneson, 2017). The assumption of axisymmetry allows the dimension of the problem to be reduced by one and hence facilitates rapid simulations.

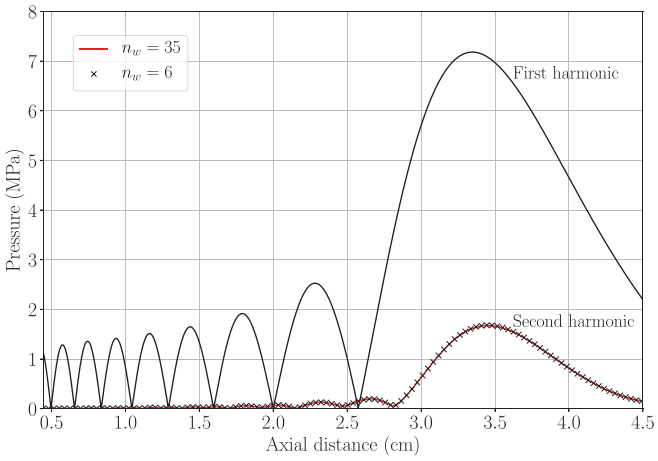

FIG. 4. (Color online) The first and second harmonics along the x axis for the H131 transducer in liver. With six voxels per wavelength, the approximation to the second harmonic is indistinguishable from the reference solution.

https://doi.org/10.1121/10.0005655

We consider two configurations:

(1) H131 transducer at output power 50 W in water;
(2) H101 transducer at output power 100 W in liver.

These power outputs are chosen because they are relevant to thermal ablation with FUS (e.g., Kothapalli *et al.*, 2018; Ries *et al.*, 2010) and are sufficiently low to ensure that we are in the weakly nonlinear setting. The first five harmonics along the x axis are shown in Figs. 5 and 6. In both cases, we observe good qualitative agreement with the approximations obtained using HITU Simulator in the region around the focus, whereas toward the transducer the two methods disagree. This is due to the parabolic assumption made in the derivation of the WAKZK equation, which is only accurate when sufficiently far from the transducer. The volume potential method on the other hand approximates solutions to the Westervelt equation and therefore can be taken to be accurate in this "near field" region. Indeed, the computation of the first harmonic in our approach, as outlined in Sec. III B, is equivalent to a Rayleigh integral method.

We observe that the volume potential method predicts a slightly larger amplitude for the first and second harmonics than HITU Simulator, which may be because the energy flow of the higher harmonics to lower harmonics is neglected in the derivation of (7). However, the agreement for the third, fourth, and fifth harmonics is almost perfect. The two examples considered have considerably different attenuation power law parameters; thus, the strong agreement with HITU Simulator for both examples demonstrates that the attenuation is being handled correctly in the volume potential method.

V. COMPUTATION DOMAINS FOR SUCCESSIVE HARMONICS

In this section, we aim to determine the amount by which we can restrict the computation domain while retaining accurate approximations and thus enable acceleration of our simulations.

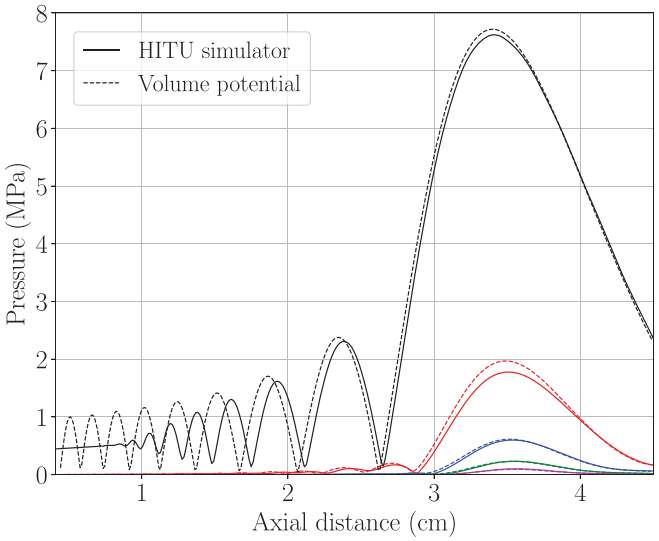

FIG. 5. (Color online) Comparison of the VIE approach with HITU Simulator for the first five harmonics generated by the H131 transducer operating at a power of 50 W in water.

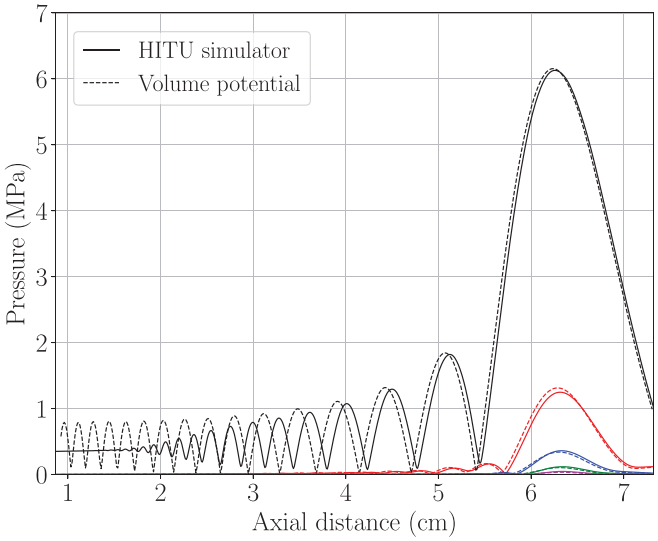

FIG. 6. (Color online) Comparison of the VIE approach with HITU Simulator for the first five harmonics. H101 transducer at power 100 W in liver.

The main region of interest to practitioners is that around the focus, since this is where tissue ablation occurs. Ideally one would compute only on that small region. This, however, leads to the second harmonic (and thus also the third, fourth, etc.) being poorly approximated, since these harmonics are generated by the accumulation of acoustic energy over distance. Thus, we seek a balance between accuracy and computational cost. Here, we shall aim to keep the error in each harmonic below 1% (relative to the magnitude of the first harmonic) whilst shrinking the integration regions D_i in (12)–(15) as much as possible.

It is important to note that (in this homogeneous setting), the field beyond the focal point has no influence on the field in front of it, i.e., the waves propagate only in the positive x-direction. This means that our exploration of domain shrinking only applies to the region between the transducer and the focal point. Beyond the focus, we keep the domain length in the x-direction fixed. In most FUS settings, the practitioners are not interested in the field far beyond the focal region; therefore, the inability to shrink the region beyond the focus does not greatly affect the gains achieved from shrinking the computational domain before the focus.

Each of Eqs. (12)–(15) has the form

$$p(\mathbf{x}) = C\int_D G_k(\mathbf{x},\mathbf{y})f(\mathbf{y})\mathrm{d}\mathbf{y}, \tag{23}$$

where $f = p_1^2, p_1p_2, \ldots$, C is the appropriate constant, and k is the appropriate wavenumber. To accurately approximate p, the integration domain D must enclose the region where the integrand is non-negligible. Outside of this region, we can discard the contributions. The magnitude of the integrand is dictated by the function f, which has the units of

intensity. To have an idea of how localised the different functions f are, we plot them for a particular example in Fig. 7. The setup considered in the figure is the H101 transducer at 100 W in liver. The figure shows the magnitudes of f scaled by their maximum values (at the focus) and converted to a log-scale. Consider the top-left image: this is the f required for the computation of the second harmonic. We can see that the magnitude is significant all the way back to the transducer, implying that we must include all this area in the integration domain D. For the remaining images, corresponding to the third, fourth, and fifth harmonics, the functions become increasingly more localised in both the x and y/z dimensions, suggesting that the required integration domain can be considerably smaller than that for the second harmonic.

To investigate this more rigorously, we perform convergence tests for each harmonic as the relevant integration domain D is restricted. That is, we take the harmonics generated on the domain in (20) as the reference solutions p_l^r, $l = 2, 3, 4, 5$, and then compute the same harmonics on successively smaller domains and compare the approximations to the reference. The error of an approximation p_l is computed along the x axis as

$$\text{Error} = \frac{||p_l - p_l^r||}{||p_1^r||} \times 100\%. \tag{24}$$

Note that the harmonic field in the denominator is that of the first harmonic. This is done so that the error function incorporates the diminishing size of successive harmonics. For example, if the fifth harmonic is negligibly small relative to the first harmonic (and so not worth calculating), the error will reflect this by being very small.

As a measure of the "localisedness" of the functions f, we use the quantity plotted in Fig. 7, which we denote as Q,

$$Q(\mathbf{x}) := \log_{10}\left(\frac{|f(\mathbf{x})|}{\max|f(\mathbf{x})|}\right). \tag{25}$$

In the convergence tests, the integration domain is chosen as the smallest cuboidal domain D such that $Q(\mathbf{x}) < Q_0$ for $\mathbf{x} \notin D$, where Q_0 is a given threshold.

The first configuration we perform the convergence experiment for is the H131 transducer in water at an output power of 100 W. The convergence for each harmonic is shown in Fig. 8. We notice that the convergence of the second harmonic drops suddenly once the error dips below 1%—this is because the computation domain is close to the size of the reference domain by this point. This is illustrated more clearly in Fig. 9, where the same data as in Fig. 8 are shown but now plotted against the size of the computation domain as a fraction of the reference domain, rather than against Q_0. By "fraction" of the domain, we mean the

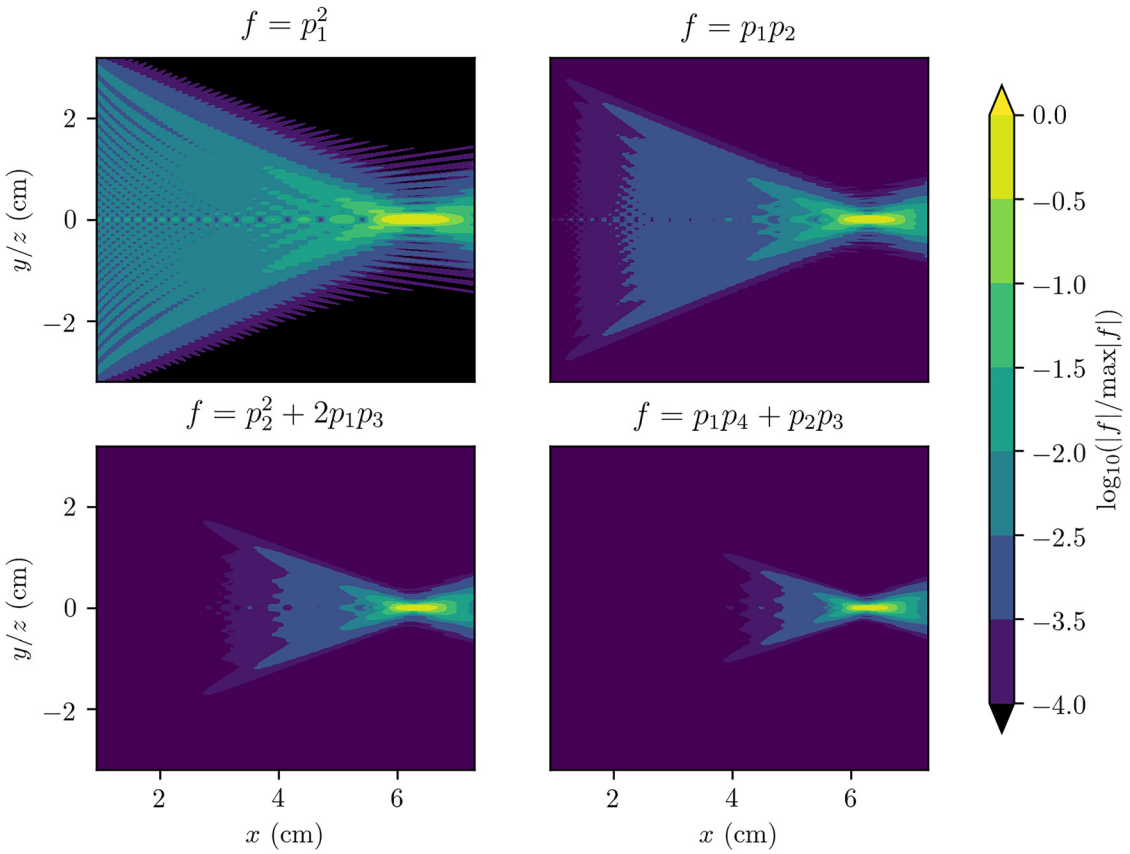

FIG. 7. (Color online) The relative magnitudes of the right-hand side functions f as described in Eq. (23) for the H101 transducer in liver, at 100 W. These plots show how the function to be convolved with the Green's function becomes more localised as the harmonics increase.

https://doi.org/10.1121/10.0005655

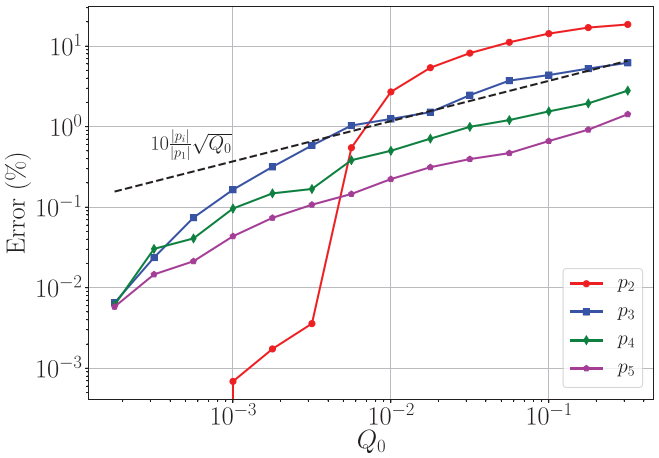

FIG. 8. (Color online) Convergence of the approximations to harmonics p_i, $i = 2, \ldots, 5$, as the domain of integration D_i for each is adjusted according to the function $Q \geq Q_0$, as defined in (25). The setup considered here is the H131 transducer at a power of 100 W in water.

scaling factor such that the length L' and width w' of the shrunken domain are given by

$$L' = \text{fraction}_x \times L, \quad w' = \text{fraction}_{y,z} \times w.$$

Figure 9 shows that, to achieve less than 1% error in p_2, the computation domain must extend all the way to the transducer. For the higher harmonics, however, a different trend is evident. In Fig. 8, we observe that the error curves for p_3, p_4, p_5 each have the approximate behaviour

$$\text{Error}(p_i) \approx 10\frac{||p_i||}{||p_1||}\sqrt{Q_0}, \quad i = 3, 4, 5, \tag{26}$$

where $||\cdot||$ represents the L^2-norm. Although an interesting observation, the utility of the relationship (26) for dictating an appropriate computation domain for p_i is not immediately apparent, since it requires the computation of $||p_i||$ before p_i has been computed. Determining an *a priori* approximation

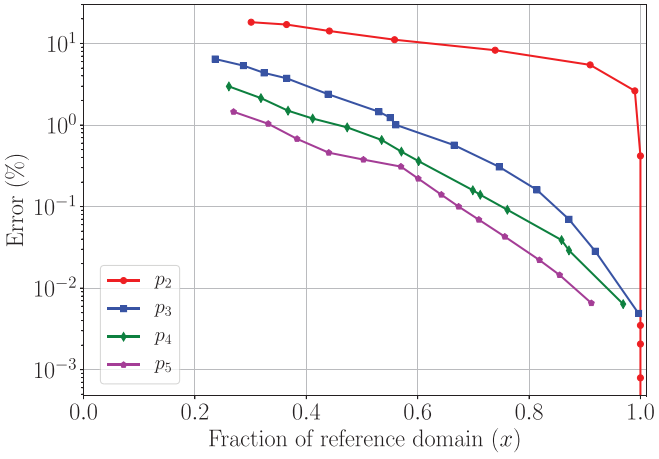

FIG. 9. (Color online) Convergence for H131 transducer at 100 W in water, as in Fig. 8. The error is plotted against the fraction of the total domain in the x-direction. This demonstrates that, to achieve an error smaller than 1%, the computation domains can be contracted significantly in the x-direction for harmonics higher than the second.

TABLE III. Sizes of domains required to achieve less than 1% error for each harmonic as fractions of the reference domain [Eq. (20)]. Since all domains are boxes, the fractions of the distances in the x and y/z directions are provided.

		p_2	p_3	p_4	p_5
H131 water, 100 W	x	1	0.67	0.47	0.38
	y/z	0.74	0.39	0.20	0.13
H131 water, 150 W	x	1	0.67	0.58	0.52
	y/z	0.71	0.40	0.39	0.21
H131 liver, 100 W	x	1	0.56	0.25	0.26
	y/z	0.90	0.20	0.05	0.06
Rule of thumb	x	1	0.75	0.65	0.61
	y/z	1	0.67	0.5	0.4

for $||p_i||$ to make use of (26) would be a useful endeavour; however, we do not undertake such a task in this article. Rather, we seek to develop an approximate rule of thumb for choosing sensibly sized computation domains for the harmonics. To this end, it is more straightforward to consider the convergence of the approximations in terms of physical distance, as is done in Fig. 9 and Table III.

By looking at the 1% error line in Fig. 9, it is possible to read off the size of the domains as fractions of the reference domain. The precise values are reported in Table III. A nested series of domains constructed according to this specification is shown in Fig. 10. Let us elaborate on the potential computational gain achieved using these nested domains compared to computing all harmonics on the reference domain.

The reference domain for the H131 transducer in water has dimensions $[4.1\,\text{cm}, 3.3\,\text{cm}, 3.3\,\text{cm}]$ and is discretised into voxels of dimension $\delta x = \lambda_5/6 \approx 45.1\,\mu\text{m}$, i.e., fine enough to resolve the highest harmonic of interest. This mesh has $901 \times 732 \times 732 \approx 4.8 \times 10^8$ voxels for the computation of all the harmonics. Using the nested domains built according to the specifications in Table III, we obtain a mesh for p_2 with dimensions $[4.1\,\text{cm}, 2.4\,\text{cm}, 2.4\,\text{cm}]$, which is discretised into voxels of dimension $\delta x = \lambda_2/6 \approx 113\,\mu\text{m}$. This mesh has

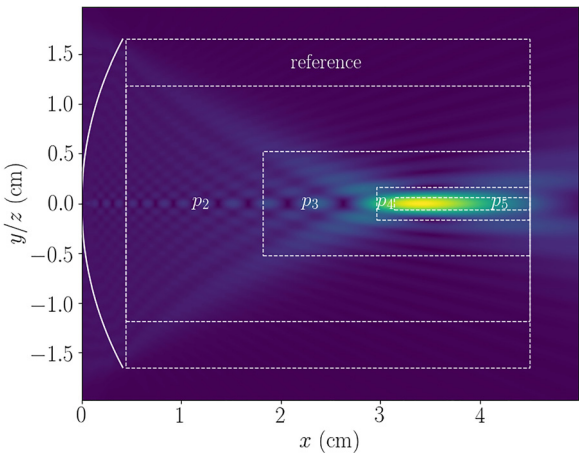

FIG. 10. (Color online) Nested domains for the computation of successive harmonics while keeping relative error below 1%, for the H131 transducer operating at 100 W in water. The domain used to compute the reference solution is in (20).

$360 \times 211 \times 211 \approx 1.6 \times 10^7$ voxels, which is a factor of 30 smaller than the reference mesh. The meshes for the computation of p_3, p_4, and p_5 have $1.22 \times 10^7, 6.1 \times 10^6$, and 4.5×10^6 voxels, respectively. Thus, the total number of voxels required for all harmonics is $4 \times 4.8 \times 10^8 = 19 \times 10^8$ for the single mesh approach and 3.9×10^7 for the nested meshing (summing over voxels in each of the four meshes). Hence, we achieve close to a factor of 50 reduction in computational load.

This improvement is impressive; however, it is easily seen in Table III that this particular domain scaling does not apply to all transducer and material configurations. Therefore, we decide upon a simple rule of thumb for designing the separate computational domains that leads to domains greater than or equal to those given in Table III. Therefore, the errors incurred in using these domains will be even smaller than those obtained when using the ideal domains specified in the table.

A desirable rule for domain sizes would be the following:

> *Let $[L + d, w, w]$ be the dimensions of domain D_2 (see Fig. 2) for harmonic p_2, with corresponding wavelength λ_2. Then choose the dimensions of domain D_i, $i > 2$, as $[(\lambda_i/\lambda_2)L + d, (\lambda_i/\lambda_2)w, (\lambda_i/\lambda_2)w]$. That is, the domain is scaled according to the wavelength of the harmonic being considered.*

With this rule for creating the domains, we observe that the number of voxels in each mesh will almost be the same, save for a slight increase due to the fact that the distance d is not being scaled. This amounts to a reduction in the overall number of DOF by a factor of approximately $(n/2)^3$ (in fact, slightly lower than this due to the unscaled portion of length d), where n is the number of harmonics being computed.

In Sec. VI, we outline an algorithm for evaluating the volume potentials over a set of nested domains constructed according to the rule of thumb proposed above.

VI. VOLUME POTENTIALS ON NESTED MESHES

To demonstrate the effectiveness of the volume potential evaluation on nested meshes, we present some final results detailing the computational performance of this approach. First, we outline the algorithm for computing the first n harmonics.

Note that in Algorithm 1, for p_2, the pressure field p_1 is evaluated over the voxel mesh $\mathcal{V}(D_2)$ as described in Sec. III B, whereas for later harmonics, we perform interpolation of the earlier harmonics down onto the new mesh. In this work, we have used linear interpolation; however, if higher accuracy is required, we recommend quadratic interpolation (albeit at a higher computational cost). The second and third step each contain applications of the FFT: for the circulant embedding of the Green's function in step 2 (see, e.g., Groth *et al.*, 2020) and to perform the convolution required in the quadrature rule (16) in step 3. These FFTs are performed using the PYTHON wrapper "pyfftw" to the FFTW library (Frigo and Johnson, 2005).

ALGORITHM 1: Algorithm for computing the first n harmonics.

for $i = 2 \rightarrow n$ **do**
1. Create domain D_i for p_i and voxel mesh $\mathcal{V}(D_i)$
2. Assemble components for integration (16):
   (a) Evaluate/interpolate $p_{i-1}, p_{i-2}, \ldots, p_1$ at voxel centres in $\mathcal{V}(D_i)$
   (b) Evaluate integral of Green's function G_{k_i} over each voxel
3. Compute p_i via appropriate equation in (12)–(15)

end for

As an example, we consider the H101 transducer operating at 100 W in water run on a workstation with two sockets, each containing a 14-core Xeon E5–2690 v4 central processing unit (CPU), each supporting hyper-threading with two threads, and hence a total number of 56 threads. The total amount of RAM available on this machine is approximately 270 GB, which is ample for the problems considered here. The first five harmonics along the x axis are shown in Fig. 11, along with the approximation obtained with HITU Simulator; again we see that the volume potential approach predicts a larger peak in the second harmonic, but in general a good qualitative agreement is observed.

To get a feel for the performance of our approach (of evaluating the volume potentials on nested meshes designed according to our proposed rule of thumb), the cost of each step in the above algorithm is detailed in Table IV. The largest mesh has 230×10^6 voxels, as compared to 2.9×10^9 voxels without nested meshing. This represents a large saving, in both time and memory. The most expensive step reported in Table IV is the evaluation of the first harmonic p_1 on $\mathcal{V}(D_2)$. This process was described in Sec. III B and has complexity $\mathcal{O}(n_e N)$, where n_e is the number of points used to discretise the surface of the transducer, and N is the

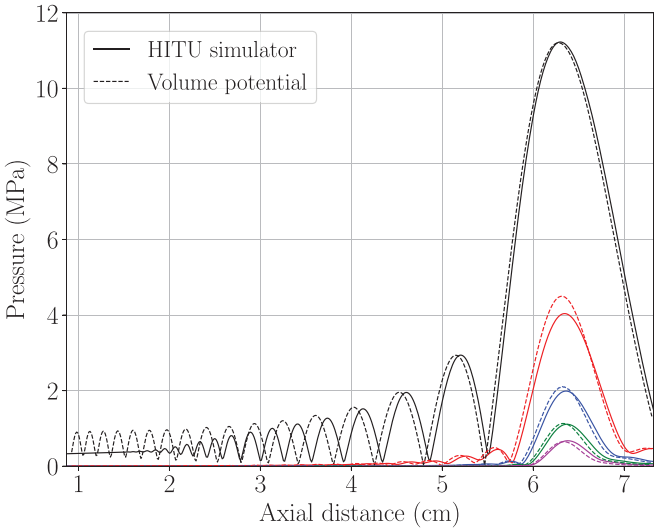

FIG. 11. (Color online) The on-axis absolute pressure field for the first five harmonics generated by the H101 transducer at 100 W in water, as computed using the volume potential approach on nested domains, designed to keep the relative error below 1%. The approximation obtained using HITU Simulator is provided for comparison.

https://doi.org/10.1121/10.0005655

TABLE IV. Performance details for the volume potential approach on nested meshes with a resolution of six voxels per wavelength (for each harmonic). The configuration considered is the H101 transducer operating at 100 W in water. Total time taken was 1 h 23 min 11 s.

	No. of voxels	Meshing	Interpolation	Evaluate G_{k_i}	Compute p_i
p_2	1.86×10^8	26.1 s	24 min 30 s	4 min 54 s	3 min 48 s
p_3	2.01×10^8	22.2 s	3 min 9 s	5 min 14 s	4 min 9 s
p_4	2.15×10^8	23.9 s	5 min 50 s	5 min 27 s	4 min 8 s
p_5	2.30×10^8	25.5 s	7 min 12 s	6 min 7 s	7 min 5 s

number of voxels. Since we take $n_e = 4096$, this is a rather expensive procedure, even when parallelised over 56 threads. Therefore, a transducer model using far fewer elements and/or a machine with more threads will lead to a large reduction in computation time for this step. Linear interpolation is computed efficiently using the SCIPY (Jones *et al.*, 2001) `RegularGridInterpolator` command. The evaluation of G_{k_i} consists of a parallelised loop over the N voxels and then an FFT of a three-dimensional complex-valued array of size $(2N_x, 2N_y, 2N_z)$, where N_x, N_y, N_z are the number of voxels in each dimension. Then the computation of p_i consists of one forward and one inverse FFT of an array of size $(2N_x, 2N_y, 2N_z)$. The FFTs take advantage of multithreading and therefore can be easily accelerated through the use of more cores. Furthermore, it seems likely that optimising the FFT routine for the particular setting and using the FFTW C++ library directly will lead to further acceleration. Nevertheless, the current implementation yields fast and accurate predictions of the first five harmonics (in the cases considered).

VII. EXTENSION TO INHOMOGENEOUS DOMAINS

We briefly consider the extension of this approach to the case of an inhomogeneous domain. That is, we suppose that the wavespeed, $c(\mathbf{x})$, and nonlinearity parameter, $\beta(\mathbf{x})$, are now spatially varying. Further, we assume that the density is close to constant, i.e., $\rho(\mathbf{x}) \approx \rho_0$ (we comment on large density contrasts at the end of this section). The spatial variation of c and β leads to backscattering of the field generated by the transducer, and thus, rather than computing the harmonics via direct evaluation of volume potentials as before, we must now in addition solve VIEs to account for the scattering effects.

Let the first harmonic of the *incident field* generated by the transducer be denoted as p^{inc}. Then VIEs for the first five harmonics are (see Costabel, 2015)

$$p_1(\mathbf{x}) - \int_{D_1} G_{k_1}(\mathbf{x},\mathbf{y})(k_1^2(\mathbf{y}) - \bar{k}_1^2)p_1(\mathbf{y})\mathrm{d}\mathbf{y} = p^{\text{inc}}(\mathbf{x}), \tag{27}$$

$$p_2(\mathbf{x}) - \int_{D_2} G_{k_2}(\mathbf{x},\mathbf{y})(k_2^2(\mathbf{y}) - \bar{k}_2^2)p_2(\mathbf{y})\mathrm{d}\mathbf{y} = -\frac{2\beta(\mathbf{x})\omega^2}{\rho_0 c(\mathbf{x})^4}\int_{D_2} G_{k_2}(\mathbf{x},\mathbf{y})p_1^2(\mathbf{y})\mathrm{d}\mathbf{y}, \tag{28}$$

$$p_2(\mathbf{x}) - \int_{D_3} G_{k_3}(\mathbf{x},\mathbf{y})(k_3^2(\mathbf{y}) - \bar{k}_3^2)p_3(\mathbf{y})\mathrm{d}\mathbf{y} = -\frac{9\beta(\mathbf{x})\omega^2}{\rho_0 c(\mathbf{x})^4}\int_{D_3} G_{k_3}(\mathbf{x},\mathbf{y})p_1(\mathbf{y})p_2(\mathbf{y})\mathrm{d}\mathbf{y}, \tag{29}$$

$$p_4(\mathbf{x}) - \int_{D_4} G_{k_4}(\mathbf{x},\mathbf{y})(k_4^2(\mathbf{y}) - \bar{k}_4^2)p_4(\mathbf{y})\mathrm{d}\mathbf{y} = -\frac{8\beta(\mathbf{x})\omega^2}{\rho_0 c(\mathbf{x})^4}\int_{D_4} G_{k_4}(\mathbf{x},\mathbf{y})(p_2^2(\mathbf{y}) + 2p_1(\mathbf{y})p_3(\mathbf{y}))\mathrm{d}\mathbf{y}, \tag{30}$$

$$p_5(\mathbf{x}) - \int_{D_5} G_{k_5}(\mathbf{x},\mathbf{y})(k_5^2(\mathbf{y}) - \bar{k}_5^2)p_5(\mathbf{y})\mathrm{d}\mathbf{y} = -\frac{25\beta(\mathbf{x})\omega^2}{\rho_0 c(\mathbf{x})^4}\int_{D_5} G_{k_5}(\mathbf{x},\mathbf{y})(p_1(\mathbf{y})p_4(\mathbf{y}) + p_2(\mathbf{y})p_3(\mathbf{y}))\mathrm{d}\mathbf{y}, \tag{31}$$

where $\bar{k}_i$ is the wavenumber of the background medium for harmonic i, and $k_i(\mathbf{x})$ is the variable wavenumber. Note that the integrals on the left-hand sides have non-zero contributions only where $k_i(\mathbf{x}) \neq \bar{k}_i$, i.e., where the wavenumber differs from that of the background medium. If $k_i(\mathbf{x}) \equiv \bar{k}_i$, then we are in the homogeneous case considered before.

To validate the rule of thumb for an inhomogeneous medium, we consider a 2 cm layer of kidney tissue surrounded by water. The layer is centred at the focus of the transducer. As the tissue properties for water and kidney, we use those given in Table II, except that we assume the density of kidney to be equal to that of water, to coincide with our constant density assumption. We consider the H131 transducer operating at 50 W. A comparison with HITU Simulator is presented in Fig. 12, where we observe good agreement in terms of focus location as well as magnitudes of the separate harmonics. Since HITU is a one-way solver, it does not approximate the backscattering, whereas our full-wave solver does. The backscattering can be observed as the ripples in the VIE curves.

We note that in the above we assumed that $\rho(\mathbf{x}) \approx \rho_0$ throughout the inhomogeneous domain. This was done to derive convenient VIEs, which can be solved in an efficient manner. For strong density contrasts, the VIEs (27)–(31) must be augmented with *boundary integrals*, as discussed in Costabel (2015). This complicated issue could be resolved by a coupling to an established boundary element code, such as in van 't Wout *et al.* (2015), but this is left to future work.

VIII. CONCLUSION

In this paper, we have set out to reduce the computational burden of numerical schemes for FUS simulations through the construction of an efficient and simple meshing strategy. This strategy can be employed with those numerical schemes that seek to approximate the Westervelt equation on a single non-uniform mesh and those that solve for each harmonic on separate meshes, such as the frequency-domain volume potential approach proposed here.

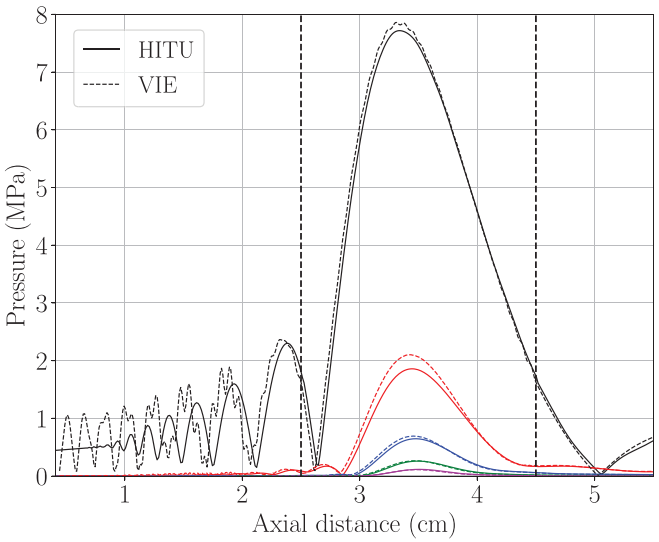

FIG. 12. (Color online) The on-axis absolute pressure field for the first five harmonics generated by the H131 transducer at 50 W in water with a 2 cm layer of kidney material centred at 3.5 cm. The vertical dashed lines demarcate the kidney layer.

The strategy exploits the increasingly localised nature of the higher harmonics around the transducer's focal region so that the DOF in the mesh can be more efficiently distributed. If we were considering a single non-uniform mesh approach in which we approximate the full Westervelt equation, this mesh would become increasingly more refined toward the focus, since this is where the higher harmonics are present. In the frequency-domain setting we considered, this leads to a nested series of meshes, as was discussed in detail in this article.

In the frequency domain, the Westervelt equation can be rewritten as a series of inhomogeneous Helmholtz equations. When the propagation medium is taken to be homogeneous, these Helmholtz equations can be solved exactly by volume potentials, which may be efficiently evaluated using the quadrature method proposed in Sec. III. This novel application of this approach allows us to explore efficiently the convergence of each harmonic as the respective computation domain was changed in size, thus enabling us to determine the smallest domains we could use to achieve an error of less than 1%.

We showed that the accurate approximation of the second harmonic requires a computation domain that extends from the focus all the way to the transducer, since the first harmonic is not sufficiently localised near the focus to allow a smaller domain to be employed. The third harmonic and above, however, can be approximated accurately on considerably reduced domains. We found that scaling the computation domain's width and height relative to the wavelength under consideration allowed for accurate approximations for the first five harmonics for the FUS configurations considered here. This leads to a reduction in the number of DOF of approximately $(n/2)^3$, where n is the number of harmonics being computed.

Finally, we demonstrated how this approach generalises, via the introduction of volume integral operators, to inhomogeneous media with low density variation. The application to inhomogeneous media with large density contrast, such as between water and bone, requires the introduction of further boundary integral operators, which is left to future work.

To conclude, we briefly comment on the generalisation of the "rule of thumb" to other transducer configurations and frequencies. In the present article, two different transducers were considered, both at 1.1 MHz and propagating within three different media: water, liver, and kidney. All the examples considered produced peak amplitudes of lower than 15 MPa, at which the weakly nonlinear assumption used in the derivation of the cascade of Helmholtz equations is accurate. We believe that for different focused transducer configurations and frequencies, our proposed rule of thumb is accurate, provided the field can still be categorised as weakly nonlinear. For highly nonlinear fields, a further study would be required to test the rule of thumb and also a modification of our volume potential approach required to allow for the transfer of energy from higher to lower harmonics.

A PYTHON implementation of this work is freely available (Groth, 2021).

ACKNOWLEDGMENTS

This work was supported by a grant entitled "Optimising patient specific treatment plans for ultrasound ablative therapies in the abdomen (OptimUS)" from the Engineering and Physical Sciences Research Council (EPSRC) (Grant Nos. EP/P013309/1 to the University of Cambridge and EP/P012434/1 to University College London).

Azhari, H. (**2010**). *Basics of Biomedical Ultrasound for Engineers* (Wiley, New York).

Bessonova, O., Khokhlova, V., Bailey, M., Canney, M., and Crum, L. (**2009**). "Focusing of high power ultrasound beams and limiting values of shock wave parameters," Acoust. Phys. **55**(4), 463–473.

Campos-Pozuelo, C., Dubus, B., and Gallego-Juárez, J. A. (**1999**). "Finite-element analysis of the nonlinear propagation of high-intensity acoustic waves," J. Acoust. Soc. Am. **106**(1), 91–101.

Chen, W., and Holm, S. (**2004**). "Fractional Laplacian time-space models for linear and nonlinear lossy media exhibiting arbitrary frequency power-law dependency," J. Acoust. Soc. Am. **115**(4), 1424–1430.

Colton, D., and Kress, R. (**2013**). Integral Equation Methods *Scattering Theory* (Society for Industrial and Applied Mathematics, Philadelphia, PA).

Costabel, M. (**2015**). "On the spectrum of volume integral operators in acoustic scattering," in *Integral Methods in Science and Engineering* (Springer, New York), pp. 119–127.

Deserno, M. (**2004**). "How to Generate Equidistributed Points on the Surface of a Sphere," Max-Planck-Institut für Polymerforschung, Mainz, Germany.

Draine, B. T., and Flatau, P. J. (**1994**). "Discrete-dipole approximation for scattering calculations," J. Opt. Soc. Am. A **11**(4), 1491–1499.

Du, Y., and Jensen, J. A. (**2013**). "Fast simulation of non-linear pulsed ultrasound fields using an angular spectrum approach," Ultrasonics **53**(2), 588–594.

Duck, F. A. (**2013**). *Physical Properties of Tissues: A Comprehensive Reference Book* (Academic, San Diego, CA).

Frigo, M., and Johnson, S. G. (**2005**). "The design and implementation of FFTW3," Proc. IEEE **93**(2), 216–231.

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

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

Groth, S. P. (**2021**). https://github.com/samuelpgroth/vines (Last viewed 7/14/2021).

Groth, S. P., Polimeridis, A. G., and White, J. K. (**2020**). "Accelerating the discrete dipole approximation via circulant preconditioning," J. Quant. Spectrosc. Radiative Transfer **240**, 106689.

Gu, J., and Jing, Y. (**2015**). "Modeling of wave propagation for medical ultrasound: A review," IEEE Trans. Ultrason. Ferroelectr. Freq. Control **62**(11), 1979–1992.

Hamilton, M. F., Blackstock, D. T., eds. (**1998**). *Nonlinear Acoustics* (Academic, San Diego, CA).

Izadifar, Z., Babyn, P., and Chapman, D. (**2017**). "Mechanical and biological effects of ultrasound: A review of present knowledge," Ultrasound Med. Biol. **43**(6), 1085–1104.

Jaros, J., Rendell, A. P., and Treeby, B. E. (**2016**). "Full-wave nonlinear ultrasound simulation on distributed clusters with applications in high-intensity focused ultrasound," Int. J. High Perform. Comput. Appl. **30**(2), 137–155.

Jones, E., Oliphant, T., and Peterson, P. (**2001**). "SciPy: Open source scientific tools for Python" http://www.scipy.org/ (Last viewed 7/14/2021).

Karzova, M. M., Yuldashev, P. V., Sapozhnikov, O. A., Khokhlova, V. A., Cunitz, B. W., Kreider, W., and Bailey, M. R. (**2017**). "Shock formation and nonlinear saturation effects in the ultrasound field of a diagnostic curvilinear probe," J. Acoust. Soc. Am. **141**(4), 2327–2337.

Khokhlova, T. D., Canney, M. S., Khokhlova, V. A., Sapozhnikov, O. A., Crum, L. A., and Bailey, M. R. (**2011**). "Controlled tissue emulsification produced by high intensity focused ultrasound shock waves and millisecond boiling," J. Acoust. Soc. Am. **130**(5), 3498–3510.

Kinsler, L. E., Frey, A. R., Coppens, A. B., and Sanders, J. V. (**1999**). *Fundamentals of Acoustics* (Wiley, New York).

Kothapalli, S. V., Partanen, A., Zhu, L., Altman, M. B., Gach, H. M., Hallahan, D. E., and Chen, H. (**2018**). "A convenient, reliable, and fast acoustic pressure field measurement method for magnetic resonance-guided high-intensity focused ultrasound systems with phased array transducers," J. Ther. Ultrasound **6**(1), 1–8.

Kreider, W., Yuldashev, P. V., Sapozhnikov, O. A., Farr, N., Partanen, A., Bailey, M. R., and Khokhlova, V. A. (**2013**). "Characterization of a multi-element clinical HIFU system using acoustic holography and nonlinear modeling," IEEE Trans. Ultrason. Ferroelectr. Freq. Control **60**(8), 1683–1698.

Marburg, S. (**2008**). "Discretization requirements: How many elements per wavelength are necessary?" in *Computational Acoustics of Noise Propagation in Fluids—Finite and Boundary Element Methods* (Springer, New York), pp. 309–332.

Maxwell, A. D., Wang, T.-Y., Cain, C. A., Fowlkes, J. B., Sapozhnikov, O. A., Bailey, M. R., and Xu, Z. (**2011**). "Cavitation clouds created by shock scattering from bubbles during histotripsy," J. Acoust. Soc. Am. **130**(4), 1888–1898.

Ries, M., De Senneville, B. D., Roujol, S., Berber, Y., Quesson, B., and Moonen, C. (**2010**). "Real-time 3D target tracking in MRI guided focused ultrasound ablations in moving tissues," Magn. Reson. Med. **64**(6), 1704–1712.

Solovchuk, M., Sheu, T. W., and Thiriet, M. (**2013**). "Simulation of nonlinear Westervelt equation for the investigation of acoustic streaming and nonlinear propagation effects," J. Acoust. Soc. Am. **134**(5), 3931–3942.

Solovchuk, M. A., Hwang, S. C., Chang, H., Thiriet, M., and Sheu, T. W. (**2014**). "Temperature elevation by HIFU in *ex vivo* porcine muscle: MRI measurement and simulation study," Med. Phys. **41**(5), 052903.

Soneson, J. (**2017**). "HITU simulator (version 2.0) [computer program]," https://www.fda.gov/about-fda/cdrh-offices/hitu-simulator (Last viewed 2020-05-06).

Soneson, J. E. (**2017**). "Extending the utility of the parabolic approximation in medical ultrasound using wide-angle diffraction modeling," IEEE Trans. Ultrason. Ferroelectr. Freq. Control **64**(4), 679–687.

Sonic Concepts (**2020**). "Sonic Concepts transducers," https://soniccon-cepts.com/transducer-selection-guide (Last viewed 2020-08-19).

Szabo, T. L. (**1994**). "Time domain wave equations for lossy media obeying a frequency power law," J. Acoust. Soc. Am. **96**(1), 491–500.

Treeby, B. E., and Cox, B. (**2010**). "Modeling power law absorption and dispersion for acoustic propagation using the fractional Laplacian," J. Acoust. Soc. Am. **127**(5), 2741–2748.

Treeby, B. E., Jaros, J., Rendell, A. P., and Cox, B. (**2012**). "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.

van 't Wout, E., Gélat, P., Betcke, T., and Arridge, S. (**2015**). "A fast boundary element method for the scattering analysis of high-intensity focused ultrasound," J. Acoust. Soc. Am. **138**(5), 2726–2737.

Vlaisavljevich, E., Kim, Y., Allen, S., Owens, G., Pelletier, S., Cain, C., Ives, K., and Xu, Z. (**2013**). "Image-guided non-invasive ultrasound liver ablation using histotripsy: Feasibility study in an *in vivo* porcine model," Ultrasound Med. Biol. **39**(8), 1398–1409.

Waters, K. R., Mobley, J., and Miller, J. G. (**2005**). "Causality-imposed (Kramers-Kronig) relationships between attenuation and dispersion," IEEE Trans. Ultrason. Ferroelectr. Freq. Control **52**(5), 822–823.

Yuldashev, P., and Khokhlova, V. (**2011**). "Simulation of three-dimensional nonlinear fields of ultrasound therapeutic arrays," Acoust. Phys. **57**(3), 334–343.

Yuldashev, P., Mezdrokhin, I., and Khokhlova, V. (**2018**). "Wide-angle parabolic approximation for modeling high-intensity fields from strongly focused ultrasound transducers," Acoust. Phys. **64**(3), 309–319.