A filtered convolution method for the computation of acoustic wave fields in very large spatiotemporal domains

Martin D. Verweij[a] and Jacob Huijssen
Laboratory of Electromagnetic Research, Faculty of Electrical Engineering, Mathematics and Computer Science, Delft University of Technology, Mekelweg 4, 2628 CD Delft, The Netherlands

(Received 22 July 2008; revised 8 January 2009; accepted 12 January 2009)

The full-wave computation of transient acoustic fields with sizes in the order of $100\times100\times100$ wavelengths by 100 periods requires a numerical method that is extremely efficient in terms of storage and computation. Iterative integral equation methods offer a good performance on these points, provided that the recurring spatiotemporal convolutions are computed with a coarse sampling and relatively few computational operations. This paper describes a method for the numerical evaluation of very large-scale, four-dimensional convolutions that employs a fast Fourier transformation and that uses a sampling rate close to or at the limit of two points per wavelength and per period. To achieve this, the functions involved are systematically filtered, windowed, and zero-padded with respect to all relevant coordinates prior to sampling. The method is developed in the context of the Neumann iterative solution of the acoustic contrast source problem for an inhomogeneous medium. The implementation of the method on a parallel computer is discussed. The obtained numerical results have a relative root mean square error of a few percent when sampling at two points per wavelength and per period. Further, the results prove that the method enables the computation of transient fields in the order of the indicated size.
 [DOI: 10.1121/1.3077220]

PACS number(s): 43.20.Px, 43.20.El, 43.20.Bi, 43.60.Gk [TDM] Pages: 1868–1878

I. INTRODUCTION

The computation of transient acoustic wave fields over very large spatiotemporal domains with inhomogeneous, dispersive, and nonlinear media plays an increasingly important role in the development and operation of novel acoustical applications and devices. For instance, consider the realization of an inverse scattering method for the accurate characterization of biomedical tissue using pulsed ultrasound data. The basis of this inverse method will be a full-wave, forward method that is repeatedly invoked for the computation of a transient acoustic pressure field in the inhomogeneous medium under reconstruction. In case of human organs and ultrasound frequencies, the forward computations may easily involve a domain that measures 100 wavelengths in three spatial directions by 100 periods in time. These forward computations provide a typical example of the kind and size of problems that have motivated the research reported in this paper.

Even with today's parallel computers, the numerical solution of full-wave, forward problems of the abovementioned size is only practically feasible if the applied numerical method is extremely efficient in terms of storage and computation. The largest classes of methods for the numerical solution of forward wave problems are finite difference methods, finite element methods, and integral equation methods. As is motivated below, in this paper iterative integral equation methods are selected for dealing with very large-scale, full-wave, forward problems. A variety of iterative methods exist for solving integral equations,[1–4] e.g., the Neumann iterative solution, over-relaxation methods, and conjugate gradient methods. All these methods aim to provide a sequence of increasingly accurate, global approximations to the wave field, and use convolutions over three-dimensional space and time to obtain these. In practice, the convolutions must be performed numerically, which requires the discretization of the computational domain with respect to its space and time coordinates. This paper describes a method for the numerical evaluation of convolutions over very large spatiotemporal domains. The discussion is limited to equidistant sampling and rectangular grids, and the symbol D_φ is used to indicate the number of grid points per wavelength and per period of a particular field component with temporal (angular) frequency φ. In this respect, the most relevant quantity is D_Φ, where Φ is the chosen maximum temporal angular frequency of interest.

With regard to storage, iterative integral equation methods require that one or more scalar field quantities should be available over the full spatiotemporal grid for further iteration. Since the total amount of grid points in the computational domain is proportional to D_Φ^4, it is paramount to make D_Φ as small as possible to avoid excessive memory requirements. However, the minimum value of D_Φ is dictated by the desired accuracy of the computations. The most memory-demanding operations in an iterative integral equation method are recurring convolutions of functions that depend on the stored iterates and kernel functions that are known in advance, e.g., the Green's function of a background medium. These convolutions involve the entire computational domain. As a consequence of the Nyquist–Shannon sampling theorem,[5] the sampling may in principle be as coarse as $D_\Phi=2$ to yield exact convolutional results at the grid points,

[a]Electronic mail: m.d.verweij@tudelft.nl

provided that components with an angular frequency above Φ are absent. For the typical example presented above, sampling at the limit $D_\Phi=2$ results in a grid of 1.6×10^9 points, and a parallel implementation of the chosen method demands a memory size in the order of 100 Gbytes. With modern, high-performance computing facilities, this is very well feasible.

From a computational point of view, the numerical convolutions may benefit from using the Fast Fourier Transformation (FFT).[6] For a convolution involving N sample points, this requires only order $N\log(N)$ operations instead of order N^2 operations for a straightforward quadrature rule. In view of this and the fact that N can be kept low by sampling close to or at the limit $D_\Phi=2$, it may be concluded that iterative integral equation methods combine a minimal storage requirement with a strongly reduced computational effort.

For comparison it is noted that standard finite difference or finite element methods use relatively simple, local functions to approximate the wave field and need at least a grid with $D_\Phi=10$ to yield accurate results.[7,8] As a consequence, for the same computational domain these methods need much more storage than an iterative integral equation method. This becomes evident when realizing that the typical example mentioned above now requires 10^{12} grid points.

Although many authors[9–12] employ FFT's for the numerical evaluation of convolutions, this is rarely accompanied by a minimization of the number of grid points. Presumably, this has to do with the fact that a coarse sampling requires special preparations, and for most applications storage is not so much of a bottleneck as computation. Since this paper deals with very large-scale wave problems, the sampling issue is at least as relevant as the application of FFT's. Therefore, this paper focuses on the development of a method for the numerical evaluation of very large-scale, four-dimensional convolutions that is based on FFT's and that uses a sampling rate close to or at the limit $D_\Phi=2$. Since filtering forms the basis of achieving the coarse sampling, this method is referred to as the filtered convolution method. The filtered convolution method is generic because it is applicable in all numerical methods that employ convolutions.

The approach to obtain a coarse sampling is to select a temporal angular cutoff frequency $\Omega\geq\Phi$ and a corresponding spatial angular cutoff frequency K, and to apply ideal filters to remove all higher temporal and spatial frequencies from the quantities to be sampled. This enables temporal and spatial samplings with $D_\Omega=2$, without the occurrence of aliasing. In addition, ideal windows are used to truncate the domain of the quantities, and zero-padding is employed. These steps are necessary to avoid wraparound due to the circular nature of the FFT-based numerical convolution.[6] The use of FFT's implies that the sampling occurs in both the original $\boldsymbol{x},t$ domain and in the transformed $\boldsymbol{k},\omega$ domain, and therefore the filtering and the windowing must in principle be performed in both domains. However, ideal filtering and ideal windowing are dual operations under the Fourier transformation, so a particular operation in one domain is automatically performed when its counterpart is performed in the other domain.

Central in this paper is a systematic explanation of the method that leads to a memory efficient and computationally fast numerical convolution. To describe this in a context, the contrast source formulation of the forward acoustic wave problem is introduced, and the Neumann iterative solution is employed to solve it. The authors are aware that other iterative integral equation methods may be more suitable for solving the problem at hand. However, selection of the best iterative method is not the issue of this paper, and the Neumann iterative solution has been chosen because it provides the most straightforward application of the filtered convolution method.

II. PROBLEM FORMULATION

A. Contrast source problem

The filtered convolution method is developed in the context of solving the system of first-order acoustic equations:

$$\boldsymbol{\nabla}p+\rho_0\partial_t\boldsymbol{v}=\boldsymbol{f}_p+\boldsymbol{f}_c(p,\boldsymbol{v}), \tag{1}$$

$$\boldsymbol{\nabla}\cdot\boldsymbol{v}+\kappa_0\partial_t p=q_p+q_c(p,\boldsymbol{v}), \tag{2}$$

in which p is the acoustic pressure and $\boldsymbol{v}$ is the particle velocity. This so-called contrast source formulation[13,14] is particularly useful for problems involving a medium with a localized inhomogeneous, anisotropic, dispersive, or nonlinear behavior. Outside the localized region, the medium behaves like a homogeneous, isotropic, dispersionless, and linear background medium with a mass density ρ_0 and a compressibility κ_0. The difference between the mass density of the actual medium and the background medium is accounted for by the contrast source density of force $\boldsymbol{f}_c(p,\boldsymbol{v})$, and the difference between the actual compressibility and the background compressibility is covered by the contrast source density of injection rate $q_c(p,\boldsymbol{v})$. The primary source that generates the acoustic wave field is represented by the volume density of force $\boldsymbol{f}_p$ and the volume density of injection rate q_p.

In many cases, combination of Eqs. (1) and (2) leads to a complete elimination of $\boldsymbol{v}$ and results in an acoustic wave equation

$$c_0^{-2}\partial_t^2 p-\boldsymbol{\nabla}^2 p=S_p+S_c(p). \tag{3}$$

Here, $c_0=(\rho_0\kappa_0)^{-1/2}$ is the wave speed of the background medium, and $S_p=\rho_0\partial_t q_p-\boldsymbol{\nabla}\cdot\boldsymbol{f}_p$ represents the primary source. The contrast source $S_c(p)$ now accounts for all differences between the actual medium and the background medium. Examples of contrast sources are

$$S_{c,\mathrm{inh}}(p)=[c_0^{-2}-c^{-2}(\boldsymbol{x})]\partial_t^2 p-\boldsymbol{\nabla}\ln[\rho(\boldsymbol{x})]\cdot\boldsymbol{\nabla}p, \tag{4}$$

$$S_{c,\mathrm{dis}}(p)=[c_0^{-2}\delta(t)-\rho(t)*_t\kappa(t)]*_t\partial_t^2 p, \tag{5}$$

$$S_{c,\mathrm{nl}}(p)=\frac{\beta}{\rho_0 c_0^4}\partial_t^2 p^2, \tag{6}$$

which, respectively, apply to an inhomogeneous medium with spatially varying $c(\boldsymbol{x})$ and $\rho(\boldsymbol{x})$, an attenuative and dispersive medium with relaxation functions $\rho(t)$ and $\kappa(t)$, and a

nonlinear medium with a coefficient of nonlinearity β. The symbol $*_t$ denotes a temporal convolution.

For the background medium, the Green's functions of the first-order acoustic equations and the acoustic wave equation are available in closed form.[15] This enables a solution of the corresponding contrast source problems by an iterative integral equation method, as shown next.

B. Neumann iterative solution

Knowing the Green's function G of the wave equation for the background medium, Eq. (3) may be recast into the integral equation[16,17]

$$p = G*_{\boldsymbol{x},t}[S_p + S_c(p)] = \mathcal{L}[S_p + S_c(p)]. \tag{7}$$

Here, the symbol $*_{\boldsymbol{x},t}$ denotes a convolution over the entire spatiotemporal computational domain, and it is assumed that the support of the primary and contrast source is located entirely within this domain. The convolution with the Green's function forms an integral operator, which is indicated by the symbol $\mathcal{L}$. Now it is supposed that the influence of the contrast source on the total field is relatively weak, i.e., that $\delta p = \mathcal{L}[S_c(p)]$ is small in comparison to $p^{(0)} = \mathcal{L}[S_p]$. Then successive approximations $p^{(j)}$ $(j=0,1,2,\ldots)$ to p may be generated by the Neumann iterative solution[3,4,18]

$$p^{(0)} = \mathcal{L}[S^{(0)}], \quad S^{(0)} = S_p, \tag{8}$$

$$p^{(j)} = p^{(0)} + \mathcal{L}[S^{(j)}], \quad S^{(j)} = S_c(p^{(j-1)}), \quad j \geq 1. \tag{9}$$

To get $\mathcal{L}[S^{(j)}] = \mathcal{L}[S^{(j)}](\boldsymbol{x},t) = \delta p^{(j)}(\boldsymbol{x},t)$, it is necessary to evaluate the four-dimensional convolution integral

$$\mathcal{L}[S^{(j)}](\boldsymbol{x},t) = \int_{\mathcal{D}_T^{(j)}} \int_{\mathcal{D}_X^{(j)}} G(\boldsymbol{x}-\boldsymbol{x}', t-t') S^{(j)}(\boldsymbol{x}',t') d\boldsymbol{x}' dt'. \tag{10}$$

For $j=0$ the integration is over the four-dimensional support $\mathcal{D}_T^{(0)} \times \mathcal{D}_X^{(0)}$ of the primary source $S^{(0)}$, and for $j \geq 1$ the integration is over the four-dimensional support $\mathcal{D}_T^{(j)} \times \mathcal{D}_X^{(j)}$ of the approximate contrast source $S^{(j)}$.

Even in the simplest cases, the convolution integral in Eq. (10) needs to be discretized and numerically evaluated. As has been discussed in Sec. I, for very large-scale problems it is vital to use a sampling that is as coarse as possible, preferably down to the limit $D_\Phi = 2$. In Secs. III and IV, it will be explained how this limit may be reached in case of a one-dimensional (temporal) convolution integral and in case of the four-dimensional (spatiotemporal) convolution integral in Eq. (10).

III. ACHIEVING $D_\Phi = 2$ IN CASE OF A ONE-DIMENSIONAL CONVOLUTION INTEGRAL

A. Basic Fourier theorems

First, it is convenient to recall two basic properties of an arbitrary function $f(t)$ and its Fourier transform

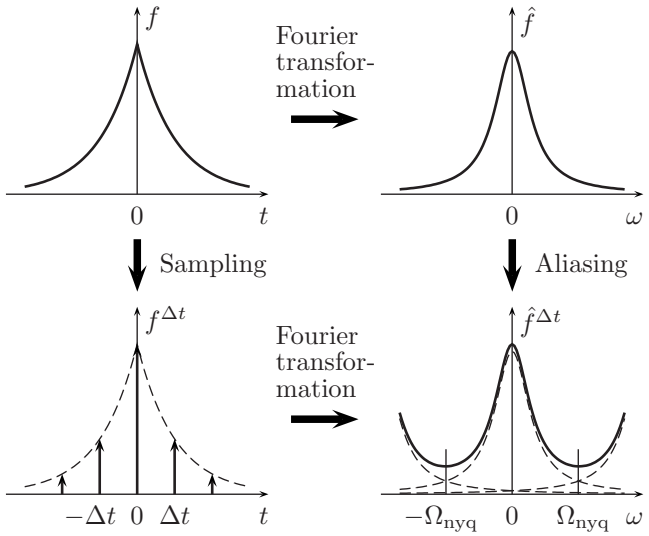

FIG. 1. Illustration of aliasing in the Fourier domain due to the sampling of a function in the original domain.

$$\hat{f}(\omega) = \int_{-\infty}^{\infty} f(t) \exp(-j\omega t) dt. \tag{11}$$

The first property involves the process of equidistant sampling of $f(t)$. If a sampled version of $f(t)$ is represented by

$$f^{\Delta t}(t) = \sum_{n=-\infty}^{\infty} f(n\Delta t)\, \delta\left(\frac{t - n\Delta t}{\Delta t}\right), \tag{12}$$

where Δt is the sampling interval, then the Fourier transform of this function is

$$\hat{f}^{\Delta t}(\omega) = \sum_{m=-\infty}^{\infty} \hat{f}(\omega + m\Omega_{\text{per}}), \tag{13}$$

which is periodic with $\Omega_{\text{per}} = 2\pi/\Delta t$. In the dual case, sampling of $\hat{f}(\omega)$ with a sampling interval $\Delta\omega$ results in a function $f^{\Delta\omega}(t)$ that is periodic with $T_{\text{per}} = 2\pi/\Delta\omega$. If $\hat{f}(\omega)$ is nonzero for $|\omega| > \Omega_{\text{per}}/2$, the terms in Eq. (13) overlap and $\hat{f}^{\Delta t}(\omega) \neq \hat{f}(\omega)$ for $|\omega| < \Omega_{\text{per}}/2$. This is the cause of aliasing and is illustrated in Fig. 1. The quantity $\Omega_{\text{per}}/2 = \pi/\Delta t = \Omega_{\text{nyq}}$ is the angular Nyquist frequency.

The second property involves the restriction of the support of $f(t)$ to an interval $[-T,T]$ according to

$$f^T(t) = f(t)[H(t+T) - H(t-T)], \tag{14}$$

where $H(t)$ is the Heaviside step function and $2T$ is the size of the applied rectangular window. The Fourier transform of this function is

$$\hat{f}^T(\omega) = \frac{T}{\pi} \hat{f}(\omega) *_\omega \operatorname{sinc}(T\omega), \tag{15}$$

in which $\operatorname{sinc}(x) = \sin(x)/x$ is the unnormalized sinc function. In the dual case, restriction of the support of $\hat{f}(\omega)$ to an interval $[-\Omega,\Omega]$ results in a function $f^\Omega(t)$ that is given by

$$f^{\Omega}(t) = \frac{\Omega}{\pi} f(t) *_t \mathrm{sinc}(\Omega t). \tag{16}$$

A finite support in the original domain always leads to an infinite support in the Fourier domain, and vice versa. In this paper, the operations to get $f^T(t)$ or $\hat{f}^T(\omega)$ are referred to as *windowing* and the operations to obtain $f^{\Omega}(t)$ or $\hat{f}^{\Omega}(\omega)$ are referred to as *filtering*.

B. Straightforward discretization of a convolution integral

Second, it is relevant to consider the numerical evaluation of a one-dimensional convolution integral

$$F(t) = G(t) *_t S(t) = \int_{-\infty}^{\infty} G(t-t')S(t')dt'. \tag{17}$$

A straightforward approach is to use the left Riemann sum with a sampling interval Δt and to approximate the convolution integral by the convolution sum[19]

$$F_n = G_n *_n S_n = \Delta t \sum_m G_{n-m} S_m. \tag{18}$$

Here, $G_n = G(n\Delta t)$ and $S_n = S(n\Delta t)$ exactly, but F_n is only an approximation of $F(t)$ at the collocation points $n\Delta t$. The exclusive use of the function values at the collocation points implies that the functions have been sampled. To keep the number of samples finite, both $n-m$ and m must be limited to a finite number, which means the windowing of $G(t)$, $S(t)$, and, consequently, $F(t)$. When the given support of S_n is $[0, N-1]$ and the desired interval of interest of F_n is also $[0, N-1]$, then the support of G_n may be restricted to $[-N+1, N-1]$ without consequences for F_n on the desired interval. For efficiency, the convolution sum is usually evaluated by applying FFT's.[20] Since this implies a circular convolution, wraparound must be avoided *a priori* by supplementing S_n with zero-valued samples S_{-N+1} to S_{-1} (zero-padding).[6] The computational effort of the FFT's thus employed is of order $2N \log(2N)$.[6,19] The remaining approximation error $F_n - F(n\Delta t)$ is due to the sampling and windowing operations. For a typical pulse $S(t)$ with a central angular frequency Ω_0, an accurate evaluation of the components of $F(t)$ up until an angular frequency of interest $\Phi = 2\Omega_0$ may easily require a sampling with $D_\Phi = 10$ or more when no prior measures are taken. This may result in a prohibitively large amount of grid points, especially for convolutions in more dimensions.

C. Coarse discretization of a convolution integral

Now the stage has been set for a discussion of the proposed method for reducing D_Φ while keeping the approximation error under control. The underlying assumption is that a maximum angular frequency of interest Φ can be chosen in such a way that all components of interest in $F(t)$ have an angular frequency $|\omega| \leq \Phi$. This implies that limitation of the spectra of $G(t)$ and $S(t)$ by filtering these functions according to Eq. (16) with an angular cutoff frequency $\Omega \geq \Phi$ only removes insignificant components of $F(t)$. By subsequently taking a sampling interval $\Delta t = \pi/\Omega$, which implies $D_\Phi = 2\pi/\Phi\Delta t = 2\Omega/\Phi$, the situation $\Omega = \Omega_{\mathrm{nyq}}$ is enforced and aliasing is avoided. This means that F_n is obtained as the exact value of $F^{\Omega}(n\Delta t)$, and the approximation error is entirely due to the difference between $F(n\Delta t)$ and $F^{\Omega}(n\Delta t)$, which manifests itself only in the part of the spectrum that is of no interest as long as $D_\Phi \geq 2$. With the proposed method, preference is thus given to discarding the contributions from the frequencies $|\omega| > \Omega$ over making an aliasing error in the frequency range $|\omega| \leq \Omega$.

So far, the approach seems ideal. However, computational reality demands a finite number of samples, and this enforces the windowing of the functions involved. Assuming that the domain of interest ranges from $t=0$ to T, a window with a size $2T$ as given in Eq. (14) is applied to both $G^{\Omega}(t)$ and $S^{\Omega}(t)$. This requires the prior formal extension of the support of $S(t)$ by defining $S(t)=0$ for $-T \leq t < 0$. The given extension corresponds to the zero-padding that is required later on. The application of the window to $G^{\Omega}(t)$ and $S^{\Omega}(t)$, in fact, re-introduces an infinite support of $\hat{G}^{\Omega,T}(\omega)$ and $\hat{S}^{\Omega,T}(\omega)$—the order of the superscripts indicating the order of corresponding operations—that will lead to aliasing. In practice, the associated error is only a fraction of the aliasing error that would occur with a direct sampling of $G(t)$ and $S(t)$.

To summarize, the filtered convolution method for the coarse discretization of a one-dimensional convolution integral involves the following:

(1) Subjecting $G(t)$ and $S(t)$ to an ideal low-pass filter with an angular cutoff frequency $\Omega \geq \Phi$, and to a time window with a size $2T = (2N-1)\pi/\Omega$.
(2) Sampling $G^{\Omega,T}(t)$ and $S^{\Omega,T}(t)$ at $2N$ points with a sampling interval $\Delta t = \pi/\Omega$.
(3) Obtaining the Fourier transforms of G_n and the zero-padded S_n using a $2N$-point FFT.
(4) Multiplying the Fourier transforms and return to the original domain using a $2N$-point inverse FFT.

This yields $F^{\Omega,T}(n\Delta t)$ as an approximation for $F(n\Delta t)$.

Alternatively, if both $\hat{G}(\omega)$ and $\hat{S}(\omega)$ are known, the dual operations may be performed in the transform domain. In this case it is convenient to first perform the windowing. Now the application of the filter to $\hat{G}^T(\omega)$ and $\hat{S}^T(\omega)$ re-introduces an infinite support of $G^{T,\Omega}(t)$ and $S^{T,\Omega}(t)$ that will give rise to time-domain aliasing. Usually, the corresponding error is only a fraction of the time-domain aliasing error that would occur with a direct sampling of $\hat{G}(\omega)$ and $\hat{S}(\omega)$.

IV. ACHIEVING $D_\Phi = 2$ IN CASE OF A FOUR-DIMENSIONAL CONVOLUTION INTEGRAL

A. Extension of the one-dimensional method

For the numerical evaluation of the four-dimensional convolution integral in Eq. (10), the method from Sec. III is applied to all four coordinates of both $G(\boldsymbol{x},t)$ and $S^{(j)}(\boldsymbol{x},t)$. The filtering and windowing operations for the spatial dimensions are denoted by the subscripts K and X, respectively.

The temporal angular cutoff frequency Ω is taken as the prime parameter for both the temporal and the spatial discretizations. As explained in Sec. III C, this quantity directly determines the temporal sampling interval $\Delta t=\pi/\Omega$. The key to the spatial sampling is the decomposition of a wave in three-dimensional space into planar wave components with real angular wavevectors $\boldsymbol{k}$, as provided by the spatial Fourier representation. For a given Ω it is now assumed sufficient to only retain wave components with $|\boldsymbol{k}|=(k_x^2+k_y^2+k_z^2)^{1/2}\leq K=\Omega/c_0$. This choice for the spatial angular cutoff frequency K is justified if the smallest details in the configuration measure at least several wavelengths at Ω. In that case the width of the spatial spectrum of the contrast source mainly depends on that of the field, and since all solutions of the source-free wave equation for the background medium have $|\boldsymbol{k}|=\omega/c_0$, the choice $K=\Omega/c_0$ seems appropriate. The intended use of FFT's implies that k_x, k_y, and k_z stay real. Consequently, in each spatial dimension the maximum possible wavenumber is K and the resulting spatial sampling interval is $\Delta x=\pi/K=\pi c_0/\Omega$. To prevent spatial aliasing with the given spatial sampling, it is *necessary* to ensure that $|k_x|\leq K$, $|k_y|\leq K$, and $|k_z|\leq K$, i.e., to apply an ideal cube-shaped filter in the spatial Fourier domain. However, as stated above, it is assumed *sufficient* to ensure that $|\boldsymbol{k}|\leq K$, i.e., to apply an ideal spherical filter in the spatial Fourier domain.

B. Spatiotemporal filtering and windowing of the four-dimensional Green's function

The four-dimensional, background Green's function occurring in Eq. (10) is[15–17]

$$G(\boldsymbol{x},t)=\frac{\delta(t-|\boldsymbol{x}|/c_0)}{4\pi|\boldsymbol{x}|}. \tag{19}$$

Its temporal Fourier domain equivalent is

$$\hat{G}(\boldsymbol{x},\omega)=\frac{\exp(-jk|\boldsymbol{x}|)}{4\pi|\boldsymbol{x}|}, \tag{20}$$

with $k=\omega/c_0$, and its spatiotemporal Fourier domain counterpart is

$$\tilde{G}(\boldsymbol{k},\omega)=\frac{1}{|\boldsymbol{k}|^2-k^2}. \tag{21}$$

It is convenient to first perform the spatial filtering in the spatiotemporal Fourier domain. Here, the spherical symmetry of $\tilde{G}(\boldsymbol{k},\omega)$ in $\boldsymbol{k}=0$ is used by applying a spherical filter with an angular cutoff frequency K in the spatial Fourier domain, giving

$$\tilde{G}_K(\boldsymbol{k},\omega)=\tilde{G}(\boldsymbol{k},\omega)H(K-|\boldsymbol{k}|). \tag{22}$$

The temporal Fourier domain equivalent of this equation has been derived by Van der Veeken and Blok[21] as

$$\begin{aligned}\hat{G}_K(\boldsymbol{x},\omega)=\hat{G}(\boldsymbol{x},\omega)&+\frac{\cos(k|\boldsymbol{x}|)}{4\pi^2|\boldsymbol{x}|}\{\mathrm{Si}[(K-k)|\boldsymbol{x}|]+\mathrm{Si}[(K+k)|\boldsymbol{x}|]\\&-\pi\}\\&+\frac{\sin(k|\boldsymbol{x}|)}{4\pi^2|\boldsymbol{x}|}\{\mathrm{Ci}[(K-k)|\boldsymbol{x}|]-\mathrm{Ci}[(K+k)|\boldsymbol{x}|]\}\end{aligned} \tag{23}$$

for $|\boldsymbol{x}|>0$, and

$$\hat{G}_K(\boldsymbol{x},\omega)=\frac{K}{2\pi^2}+\frac{k}{4\pi^2}\left[\ln\left(\frac{K-k}{K+k}\right)-j\pi\right] \tag{24}$$

for $|\boldsymbol{x}|=0$. In Eq. (23), $\mathrm{Si}(x)$ and $\mathrm{Ci}(x)$ are the sine and cosine integrals.[22]

Next, the temporal windowing is performed in the temporal Fourier domain. For the *unfiltered* Green's function, this yields

$$\begin{aligned}G^T(\boldsymbol{x},t)&=\frac{\delta(t-|\boldsymbol{x}|/c_0)}{4\pi|\boldsymbol{x}|}[H(t+T)-H(t-T)]\\&=\frac{\delta(t-|\boldsymbol{x}|/c_0)}{4\pi|\boldsymbol{x}|}H(T-|\boldsymbol{x}|/c_0),\end{aligned} \tag{25}$$

and therefore

$$\hat{G}^T(\boldsymbol{x},\omega)=\frac{\exp(-jk|\boldsymbol{x}|)}{4\pi|\boldsymbol{x}|}H(T-|\boldsymbol{x}|/c_0). \tag{26}$$

This result can be applied directly to the first term of Eq. (23), i.e., the unfiltered Green's function. The time windowing of the other terms in Eq. (23), which represent the spatial filtering operation, is more involving. When the time windowing of these terms is simply omitted, a small time-domain aliasing error is made during the evaluation of the temporal convolution sum. As explained at the end of Sec. III C, the temporal filtering of the term $G^T(\boldsymbol{x},\omega)$ that must still take place is going to yield a comparable time-domain aliasing error, so it is decided to restrict the time windowing to the first term of Eq. (23) and to leave the other terms in this equation unaffected. The remaining temporal filtering and spatial windowing operations are performed in the temporal Fourier domain as well by multiplying $\hat{G}_K^T$ with the appropriate rectangular windows. This results in $\hat{G}_{K,X}^{T,\Omega}(\boldsymbol{x},\omega)$.

C. Spatiotemporal filtering of the primary sources

In all practical cases, the spatial support of the primary source $S^{(0)}(\boldsymbol{x},t)$ is limited and the spatial windowing is unnecessary. Likewise, pulsed sources usually have a limited temporal support that makes the temporal windowing superfluous.

When the source signature has a temporal spectrum that is insignificant beyond Ω, the temporal filtering may be skipped as well without the danger of introducing noticeable aliasing. In many other situations, the Fourier transform of the signature is known in analytical form, so the temporal filtering may conveniently be performed in the temporal Fourier domain. The spatial filtering may be applied in a

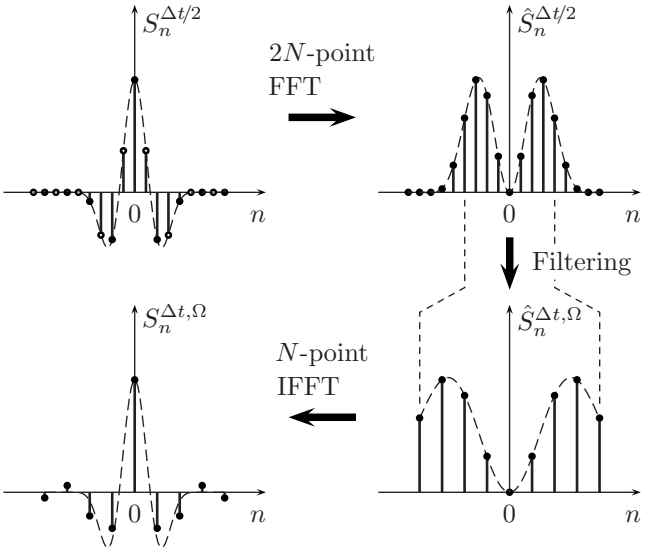

FIG. 2. Numerical procedure for filtering. In this example, the intended number of gridpoints is $N=9$ and the oversampling factor is $a=2$. In the upper left graph, the intermediate samples are indicated by open circles. The dashed curves apply to the unfiltered continuous function.

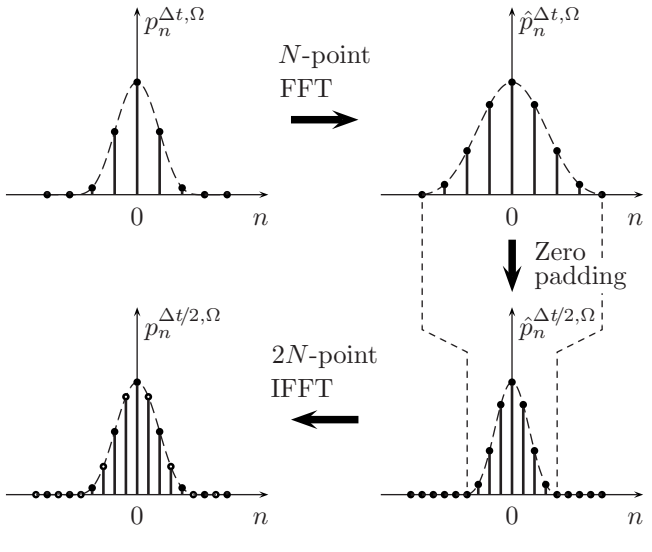

FIG. 3. Numerical procedure for interpolation. In this example, the available number of gridpoints is $N=9$ and the oversampling factor is $a=2$. In the lower left graph, the intermediate samples are indicated by open circles. The dashed curves apply to the continuous function.

similar way whenever the Fourier transform of the source geometry is available in analytical form, e.g., for rectangular or cylindrical piston sources, or for phased array transducers.[23] In such cases, the sampling is directly applied in the Fourier domain. Analytical spatial filtering avoids problems such as staircase effects or the missing out of small sources that are located in between the gridpoints.

When for some coordinate the filtering of the source cannot be employed before the sampling, e.g., when for the relevant coordinate the source can only be characterized through discrete measurements, aliasing must be avoided in a numerical way. For this purpose, the procedure as exemplified in Fig. 2 is proposed. The procedure starts with oversampling the relevant source window at aN points instead of the intended N points. Subjecting the samples to an aN-point FFT yields a discrete version of the Fourier domain function over an interval that is a times as long as intended but, for a properly chosen oversampling factor a, with negligible aliasing at the N points around the origin. Next, all points except those N points around the origin are discarded, thus giving the Fourier domain function over an interval of the intended length. This restriction is the actual filtering step since the remaining discrete interval corresponds with the continuous angular frequency interval $[-\Omega,\Omega]$. Finally, an N-point inverse FFT is performed. This yields N samples in the original domain, at the intended sampling interval and without aliasing. The proposed approach is suitable for the numerical filtering with respect to the spatial coordinates as well.

D. Spatiotemporal filtering of the contrast source

The specific behavior of the contrast source $S^{(j)}(\boldsymbol{x},t)$ depends on the type of contrast present in the medium, as shown at the end of Sec. II A. In any case, $S^{(j)}$ depends on the estimate $p^{(j-1)}$. To discuss the filtering of the contrast source, it is assumed that $p^{(j-1)}$ is bandlimited in all dimensions, with a temporal angular cutoff frequency Ω and a spatial angular cutoff frequency K. Differentiation or convolution of a bandlimited signal yields a result with a bandwidth not exceeding that of the original signal. However, multiplication of a bandlimited signal with another bandlimited signal or with itself causes a result with a bandwidth that is equal to the sum of the bandwidths of both signals involved. Thus, in case of dispersion [Eq. (5)], the contrast source does not require additional filtering. But in case of inhomogeneity [Eq. (4)], the contrast source must be filtered in the spatial dimensions, and in case of nonlinearity [Eq. (6)], the source must be filtered in both the spatial and the temporal dimensions. Detailed analysis shows that for nonlinearity the temporal filtering usually suffices and the spatial filtering may be skipped, but this issue is beyond the scope of this paper.

The abovementioned facts are true for continuous functions as well as their sampled counterparts. Usually, only the sampled version of $p^{(j-1)}$ is available and $S^{(j)}$ must be determined numerically. When this requires the sampled version of a contrast function, e.g., $[c_0^{-2}-c^{-2}(\boldsymbol{x})]$, this must be filtered either analytically (before the sampling) or numerically (after the sampling) to avoid aliasing. When the sampled function $p^{(j-1)}$ is multiplied with a sampled contrast function or with itself, the original sampling intervals $\Delta t=\pi/\Omega$ and $\Delta x=\pi/K$ are no longer sufficient to deal with the increased bandwidth of the result. To avoid aliasing, the relevant sampling intervals must be reduced by some factor a before the multiplication or the squaring takes place. This requires the interpolation of the available samples of $p^{(j-1)}$ without the distortion of the corresponding spectral values. To achieve this, the procedure as depicted in Fig. 3 is suggested. First, an N-point FFT is applied to the available N sample points. The discrete Fourier domain is then extended from N to aN points by zero-padding. Finally, an aN-point inverse FFT is performed. This yields aN samples in the original domain, without distorting the relevant spectral values.

Using the interpolated samples of $p^{(j-1)}$, the relevant

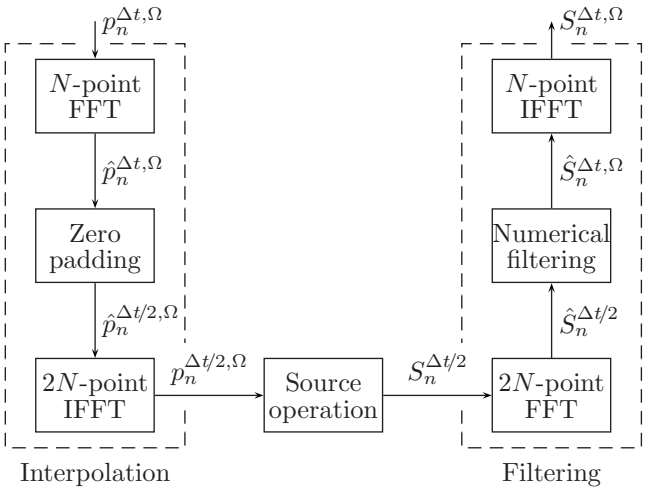

FIG. 4. The entire numerical procedure for the evaluation of the contrast source.

source operation yields aN samples of $S^{(j)}$. This oversampled $S^{(j)}$ is directly amenable for a subsequent numerical filtering according to Fig. 2. The entire numerical procedure for the evaluation of the contrast source is summarized in Fig. 4. At this point, the evaluation of the filtered contrast source may seem computationally involving. However, the contrast source very seldom requires filtering in all four dimensions. When filtering is only required for a subset of the dimensions, the number of operations required for the filtering is at least one order less than for the four-dimensional convolution, and it always pays to reduce the number of grid points by using a proper filtering.

E. Discretization and evaluation of the convolution sum

The filtered and windowed Green's function from Sec. IV B allows for a straightforward sampling. In Secs. IV C and IV D, it has been explained how sampled versions of the filtered and windowed primary source and contrast source are obtained. Apart from the small error introduced by combining filtering and windowing, all aliasing errors have been removed and therefore the sampling may be taken as coarse as $D_\Phi=2$ in all dimensions. The resulting convolution sum can be evaluated efficiently by using a four-dimensional FFT method. To avoid wraparound, all source intervals must be doubled in length by zero-padding, as explained in Sec. III B. As a result, the number of operations required for the convolution is of order $16N^4 \log(16N^4)$.

V. PARALLEL IMPLEMENTATION

The filtered convolution method has been implemented on a single processor system as well as on a multiprocessor parallel system with distributed memory. In both cases, the code has been written in FORTRAN. This section focuses on the parallel implementation[24] of the method as it occurs in the context of the Neumann iterative solution. To ease the notation, the filtering and the windowing of the acoustic quantities are understood without specific indication.

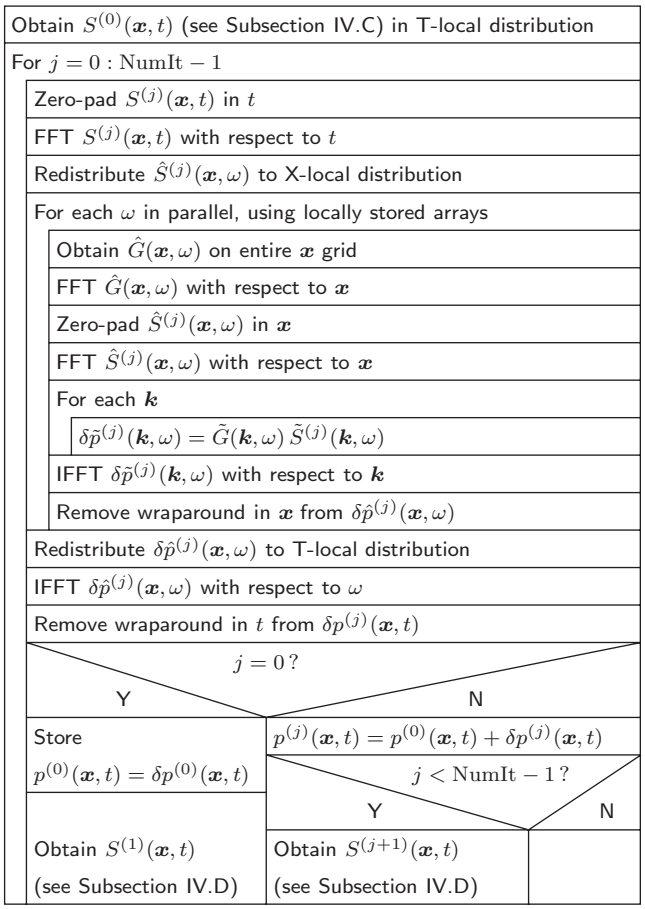

FIG. 5. Structure diagram of the parallel program. All symbols indicating acoustic quantities refer to the filtered and windowed quantities. Note that $\delta p^{(j)}=\mathcal{L}[S^{(j)}]$.

The structure diagram of the parallel program is shown in Fig. 5, in which the evaluation of the convolution plays a central role. To minimize communication, depending on the operation to be applied, either a T-local distribution is used, where each processor is allocated a complete time trace for a limited number of spatial grid points, or an X-local distribution is used, where each processor is allocated the entire spatial domain for a number of time instances or angular frequencies. Inside the outer loop over j, the source function $S^{(j)}(\boldsymbol{x},t)$ is zero-padded in time, transformed in the temporal dimension, and redistributed to X-local. This way, inside the nested loop over ω, the local temporary arrays that hold the spatial data including the zero-padding regions need only be three-dimensional. Another advantage of this approach is that the redistribution steps occur outside the loop over ω. Inside this loop, $\hat{G}(\boldsymbol{x},\omega)$ is evaluated for each spatial point. Either a high-order finite difference scheme or a spectral derivative[25] is used for the evaluation of the temporal derivatives that occur in the primary source function $S^{(0)}(\boldsymbol{x},t)$ and in the contrast source function $S^{(j)}(\boldsymbol{x},t)$ (see Sec. II A). Especially when the numerical interpolation of Fig. 3 is invoked, the route via the transform domain yields the spectral derivative at negligible extra cost. When this option is chosen, care is taken to avoid Gibb's phenomenon[20] that might arise from a jump between the data at the beginning and at the end of the interval that will be transformed by the FFT.

As the filtered convolution method relies heavily on the use of multi-dimensional FFT's, a fast and flexible implementation of the FFT is essential. This is achieved by using the FFTW library.[26]

VI. NUMERICAL RESULTS

The results in this section have been obtained by running the parallel program on an SGI Altix 3700 clustered multiprocessor system of Itanium 2 processors running at 1.3 GHz and with 2 Gbyte memory per processor.

A. The primary source

In the examples that follow, the primary source is flat and located in the plane $z=0$. The source function either follows from a volume density of injection rate $q=Q(x,y,t)\delta(z)$, giving $S=\rho_0\partial_t Q(x,y,t)\delta(z)$, or from a volume density of force $f_z=F_z(x,y,t)\delta(z)$, giving $S=-\partial_z[F_z(x,y,t)\delta(z)]$. The action of these respective sources will be represented by a jump Δv_z in the z component of the particle velocity or a jump Δp in the acoustic pressure. These jumps are only nonzero at the location of the source aperture. For the emitted wave field, symmetry properties make that for $z>0$ these source/saltus descriptions relate to the boundary value descriptions of a source in a perfectly rigid baffle and in a perfectly compliant baffle, respectively, by taking $Q(x,y,t)=2v_{z0}(x,y,t)$ or $F_z(x,y,t)=2p_0(x,y,t)$, where v_{z0} and p_0 are the normal velocity and the pressure prescribed on the boundary. The source functions $Q(x,y,t)$ and $F_z(x,y,t)$ are separated into a time signature and a geometry factor $0\leq A(x,y)\leq 1$. For the time signature, a harmonic wave modulated by a Gaussian envelope is taken. These choices give

$$\begin{pmatrix} Q(x,y,t) \\ F_z(x,y,t) \end{pmatrix} = \begin{pmatrix} Q_0 \\ F_0 \end{pmatrix} A(x,y)\exp[-(2t/t_w)^2]\sin(\omega_0 t), \quad (27)$$

where Q_0 and F_0 are the peak source values, t_w is the pulse width, and ω_0 is the angular center frequency. If there are any position-dependent delays of the time signature, as with focused sources, these may be accounted for in the temporal Fourier domain by using a complex geometry factor

$$\hat{A}(x,y,\omega) = A(x,y)\exp[-j\omega t_d(x,y)], \quad (28)$$

where $t_d(x,y)$ is the delay time. With regard to the temporal filtering of S, the pulse width t_w is taken such that the significant part of the spectrum always occurs below the cutoff frequency, and the filtering operation is omitted. Spatial filtering in the x and y dimensions is performed with the numerical method of Fig. 2. In the z dimension, filtering of the delta function with an ideal filter gives

$$\frac{K}{\pi}\delta(z) *_z \operatorname{sinc}(Kz) = \frac{1}{\Delta x}\operatorname{sinc}(\pi z/\Delta x), \quad (29)$$

which for $z=n\Delta x$ is $1/\Delta x$ for $n=0$ and 0 for $n\neq 0$. In case of a force source, the spatial derivative with respect to z is not directly applied to the delta function but to the result of the convolution, according to the property $G*_z(\partial_z f_z)=\partial_z(G*_z f_z)$.

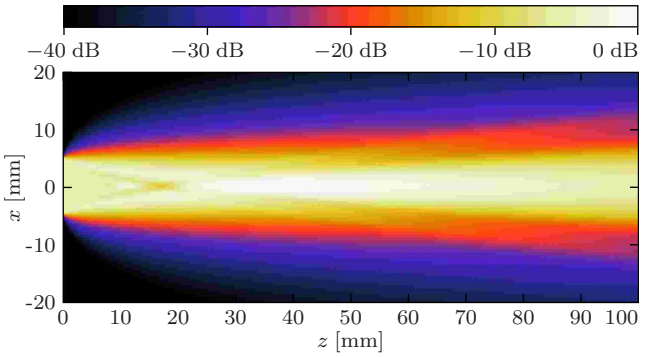

FIG. 6. (Color online) Pressure profile in the plane $y=0$ mm for the cylindrical piston transducer in water. The 0 dB level corresponds to 19.4 kPa. The sampling is $D_F=2.33$.

B. Cylindrical piston transducer in a homogeneous medium

The first example involves a cylindrical piston transducer with a radius of 5 mm, exciting a pulse with a center frequency $f_0=2$ MHz, a pulse width $t_w=3/f_0$, and a maximum surface pressure $p_0=10$ kPa. The action of the transducer is represented by a pressure jump condition. The transducer is spatially filtered by oversampling with a factor $a=8$ followed by numerical filtering, and the derivative in the source function is evaluated with a 30th-order finite difference scheme. The acoustic medium is water with $\rho_0=998$ kg m^{-3} and $c_0=1480$ m s^{-1}. For the applied pressures and frequencies, this medium may be considered linear and lossless. The homogeneous field solution $p^{(0)}(\boldsymbol{x},t)$ has been determined in a computational domain with an approximate spatial size of $40\times 2\times 100$ mm^3 and a co-moving time frame of duration $20/f_0$. The discretization of this domain is given by $D_F=1/F\Delta t$, where the maximum frequency of interest $F=\Phi/2\pi$ is taken 3 MHz. At this frequency, the spectrum of the excitation pulse is 50 dB below the level of the center frequency. Figure 6 shows the computed pressure profile, i.e., the maximum temporal acoustic pressure, in the plane $y=0$ mm for $D_F=2.33$.

To assess the accuracy of the results, in the plane $y=0$ mm the relative root mean square (RRMS) error with respect to a reference solution p_{ref} is considered. This reference follows from a Lobatto integration of an exact expression for the field of the pressure jump source, as derived with the coordinate transformation method described by Harris.[27] The RRMS error is given by

$$\text{Err} = \left(\int_{\mathcal{V}} (p^{(0)} - p_{\text{ref}})^2 dxdzdt \Big/ \int_{\mathcal{V}} (p^{(0)})^2 dxdzdt \right)^{1/2}, \quad (30)$$

where $\mathcal{V}$ is the total xzt-domain. The integrals are evaluated using the left Riemann sum. Table I presents the RRMS error and some other information for different values of D_F. These numerical experiments show that the results for $D_F=2.33$ compare quantitatively very well with the reference solution. The arrival times of the body and edge waves are reproduced correctly. This is also the case for $D_F=2.00$, although for this discretization the error is 4%. The primary cause for this

TABLE I. Data for several computations involving the cylindrical piston transducer in water. The computations are performed for different values of D_F.

D_F	Err (%)	4D grid size ($\times 10^6$ points)	Wall clock time (s)	Number of processors
2.00	4.0	25.72	143	4
2.33	1.0	39.26	230	4
2.67	0.8	60.24	393	4
3.33	0.5	155.7	873	4
4.00	0.5	345.6	686	8
5.33	0.4	1150	1203	48

error is the high-order finite difference scheme that is used for the spatial derivative in the source function. For higher values of D_F, the error decreases quickly to 0.4% at $D_F=5.33$. The main contribution to the latter error comes from the spatial filtering, which is most prominent at the first plane beyond the source. If the results for this plane are left out, the errors for D_F being equal to 3.33, 4.00, and 5.33 decrease to 0.2%, 0.1%, and 0.1%, respectively.

The example proves that domains up to approximately 500×10^6 grid points can be handled successfully and with a limited number of processors. This number of grid points corresponds to a four-dimensional computational domain of 75 wavelengths/periods in each dimension. The latter observation clearly indicates the applicability of the filtered convolution method to very large-scale problems.

As can be seen from the last line of Table I, the problem size and computation time grow massively for increasing values of D_F. This illustrates once more the necessity of using a coarse discretization.

C. Phased array transducer in a homogeneous medium

The second example concerns a 64 element phased array transducer with elements measuring 0.25 mm (width in the x direction) by 12 mm (height in the y direction), and a pitch of 0.3 mm. The elements excite the same surface pressure as the cylindrical transducer mentioned above, and the array is focused at $z=60$ mm, including elevation focusing. The results of the current method are compared with the results from the FIELDII program.[28] Since this program solves the Rayleigh integral,[27] the action of the transducer is now represented by a velocity jump. The phased array source is spatially filtered in the x dimension by analytically filtering one element and subsequently stacking the results, and in the y dimension by numerically filtering a single element by the method of Fig. 2. The temporal derivative in the source function is performed analytically. Again, the medium is water. The homogeneous field solution $p^{(0)}(\boldsymbol{x},t)$ has been determined in a computational domain with an approximate spatial size of $80\times 1\times 100$ mm^3 and a co-moving time frame of duration $40/f_0$. Figure 7 shows the computed pressure profile in the plane $y=0$ mm for $D_F=2.00$ with $F=3$ MHz.

Table II presents the RRMS error in the plane $y=0$ mm, together with other data relating to several discretizations. The listed error has been determined relative to the solution of FIELDII, which was run with a sampling frequency of 400 MHz and 20×40 mathematical elements per physical element. Table II shows that already for $D_F=2.00$ the results from both programs compare very well. This clearly proves that the filtered convolution method is capable of accurately computing the space-time domain acoustic field of complex plane sources, even when sampled at the limit $D_F=2$.

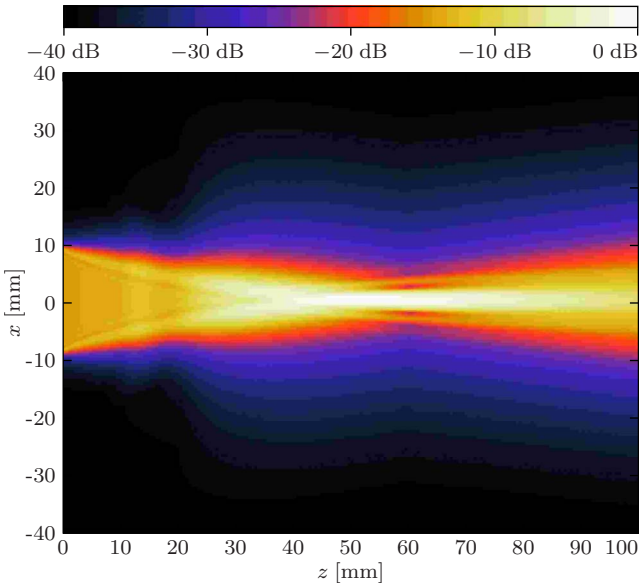

FIG. 7. (Color online) Pressure profile in the plane $y=0$ mm for the phased array transducer in water. The 0 dB level corresponds to 47 kPa. The sampling is $D_F=2.00$.

D. Phased array transducer in an inhomogeneous medium

The third example concerns a phased array transducer with the same geometry and surface pressure as in Sec. VI C. This transducer, however, excites an acoustic field in a medium with an inhomogeneous wave speed and a constant mass density. The distribution of the wave speed in the medium is displayed in Fig. 8(a). The homogeneous embedding is water. This surrounds a homogeneous phase screen ($c=0.95c_0$) with an undulation in the x direction, a water-filled cylinder ($c=0.95c_0$) with a homogeneous wall, and a prolate ellipsoid with an aspect ratio of 2 and a continuously changing wave speed ($\max[c]=1.10c_0$). The inhomogeneous wave speed leads to the contrast source of Eq. (4). The contrast function $[c_0^{-2}-c^{-2}(\boldsymbol{x})]$ is filtered by using the method of Fig. 2 with an oversampling factor $a=2$. Successive approximations $p^{(j)}$ up to $j=9$ have been determined in a computational domain with a spatial size of $42\times 16\times 60$ mm^3 and a

TABLE II. Data for several computations involving the phased array transducer in water. The computations are performed for different values of D_F.

D_F	Err (%)	4D grid size ($\times 10^6$ points)	Wall clock time (s)	Number of processors
2.00	1.6	67.0	352	4
2.67	1.0	161.6	510	8
3.33	0.7	318.1	1380	8

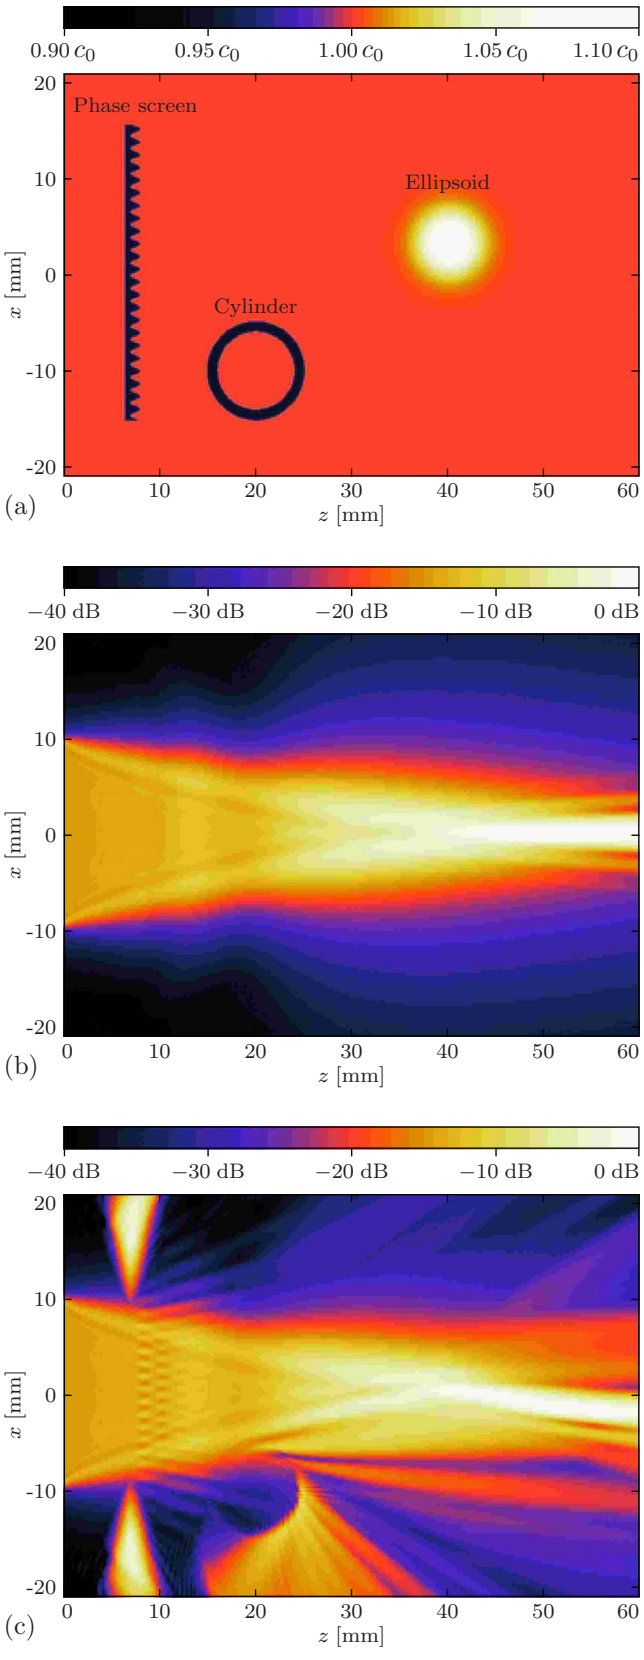

FIG. 8. (Color online) The relative wave speed (a), the pressure profile of $p^{(0)}$ (b), and the pressure profile of $p^{(9)}$ (c) in the inhomogeneous medium, for the plane y=0 mm. In panels (b) and (c), the 0 dB level corresponds to 47 kPa and the sampling is D_F=2.66.

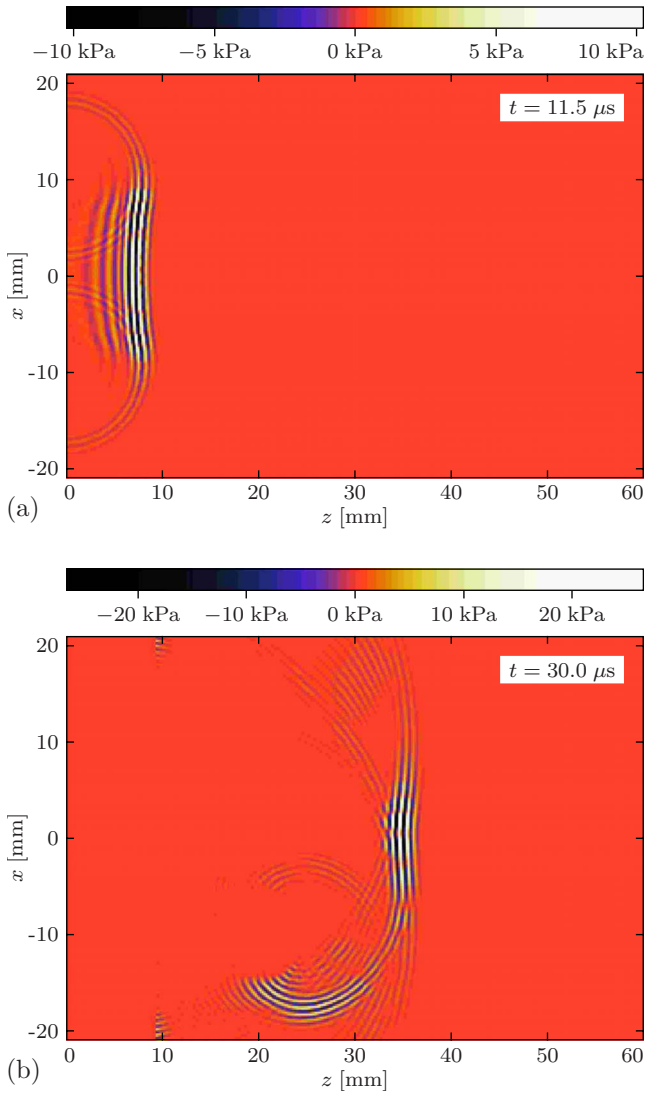

FIG. 9. (Color online) Snapshots of $p^{(9)}$ in the inhomogeneous medium, for the plane y=0 mm. The sampling is D_F=2.66.

co-moving time frame of duration $48/f_0$. Figures 8(b) and 8(c) show the profiles of $p^{(0)}$, i.e., the incident field, and $p^{(9)}$, i.e., an approximation of the total field, in the plane y=0 mm.

The results in Fig. 8(c) clearly reveal the effect of the objects on the beam shape. The phase screen affects the beam shape by generating an interference pattern just beyond the screen and by forming two side-lobes. The wall of the cylinder acts as a curved waveguide, which captures a part of the beam and radiates it off in the negative x direction. The main effect of the ellipsoid, which acts as a diverging acoustic lens, is the diffraction of the beam in the negative x direction. In Fig. 9, which shows snapshots of $p^{(9)}$, the reflections and interferences due to the inhomogeneities are also clearly visible.

In the current case, the successive results are almost identical from j=5 to j=9. However, due to the lack of benchmarking results, it has neither been investigated whether in the given situation the Neumann iterative solution is convergent, nor whether a possible convergence is toward the correct result.[29] In any case, the filtered convolution

method may be combined with more sophisticated iterative schemes[3] with improved convergence properties, such as the successive over-relaxation method or the conjugate gradient method.

The results above have been obtained for $D_F=2.66$ with $F=3$ MHz. For the computation of the successive approximations, the entire spatial domain of $85\times32\times121$ wavelengths and a co-moving time interval of 72 periods, all relating to F, acted as the support of the contrast source. The spatiotemporal grid consisted of 1.16×10^9 points, and it took 4.3 wall clock hours with 48 processors to obtain the iterations up to $j=9$. This proves that very large-scale, complex acoustic wave fields can efficiently be computed with the filtered convolution method.

VII. CONCLUSIONS

A method has been developed for the numerical evaluation of convolutions over very large, four-dimensional domains. The method is quite general and may, for instance, be combined with various iterative integral equation methods for the solution of numerous contrast source problems. In this role, the presented method enables the computation of very large spatiotemporal wave fields by minimizing the storage requirement and strongly reducing the computational effort.

As commonly encountered, the numerical convolutions are performed with the aid of FFT's. The distinguishing feature of the method is that before sampling, the relevant functions are systematically filtered, windowed, and zero-padded with respect to all relevant coordinates. Both analytical and numerical filtering procedures have been presented. Due to these preparatory steps, it is possible to perform numerical convolutions with a sampling that approaches the limit $D_\Phi=2$, i.e., two points per wavelength and per period of the highest angular frequency of interest Φ.

The method has been combined with the Neumann iterative solution of the acoustic contrast source problem for an inhomogeneous medium. Within this context, the method has been implemented on a parallel computer, and numerical results have been obtained for several configurations. Comparison with two separate reference solutions for a homogeneous medium shows that the RRMS error of our results is a few percent for $D_\Phi=2$, and decreases rapidly for finer sampling. The major part of this error is due to the filtering operation, and is thus made consciously. Computation of the transient acoustic wave field in an inhomogeneous medium proves that with the presented method, computational domains in the order of 100 wavelengths in three spatial directions by 100 periods in time can conveniently be dealt with.

ACKNOWLEDGMENTS

This research was supported by the Dutch Technology Foundation (STW) and the Dutch National Computing Facilities Foundation (NCF).

[1]P. van den Berg, "Iterative computational techniques in scattering based upon the integrated square error criterion," IEEE Trans. Antennas Propag. **32**, 1063–1071 (1984).
[2]P. van den Berg, in *PIER 5: Application of Conjugate Gradient Method to Electromagnetics and Signal Analysis*, edited by T. Sarkar (Elsevier, New York, 1991), pp. 27–65.
[3]R. Kleinman and P. van den Berg, in *PIER 5: Application of Conjugate Gradient Method to Electromagnetics and Signal Analysis*, edited by T. Sarkar (Elsevier, New York, 1991), pp. 67–102.
[4]R. Kleinman and P. van den Berg, "Iterative methods for solving integral equations," Radio Sci. **26**, 175–181 (1991).
[5]C. Shannon, "Communication in the presence of noise," Proc. IEEE **86**, 447–457 (1998).
[6]W. Press, S. Teukolsky, W. Vetterling, and B. Flannery, *Numerical Recipes in Fortran 77: The Art of Scientific Computing*, 2nd ed. (Cambridge University Press, Cambridge, 1996).
[7]A. Taflove, *Computational Electrodynamics: The Finite Difference Time-Domain Method* (Artech House, Boston, 1995).
[8]J. Jin, *The Finite Element Method in Electromagnetics*, 2nd ed. (Wiley, New York, 2002).
[9]N. Bojarski, "The *k*-space formulation of the scattering problem in the time domain," J. Acoust. Soc. Am. **72**, 570–584 (1982).
[10]Q. Liu, "Generalization of the *k*-space formulation to elastodynamic scattering problems," J. Acoust. Soc. Am. **97**, 1373–1379 (1995).
[11]T. Mast, L. Souriau, D.-L. Liu, M. Tabei, A. Nachmann, and R. Waag, "A *k*-space method for large-scale models of wave propagation in tissue," IEEE Trans. Ultrason. Ferroelectr. Freq. Control **48**, 341–354 (2001).
[12]M. Tabei, T. Mast, and R. Waag, "A *k*-space method for coupled first-order acoustic propagation equations," J. Acoust. Soc. Am. **111**, 53–63 (2002).
[13]P. Morse and K. Ingard, *Theoretical Acoustics* (McGraw-Hill, New York, 1968).
[14]A. de Hoop, *Handbook of Radiation and Scattering of Waves* (Academic, San Diego, 1995).
[15]L. Felsen and N. Marcuvitz, *Radiation and Scattering of Waves* (IEEE, New York, 1994).
[16]G. Barton, *Elements of Green's Functions and Propagation: Potentials, Diffusion and Waves* (Oxford University Press, Oxford, 1989).
[17]J. DeSanto, *Scalar Wave Theory: Green's Functions and Applications* (Springer-Verlag, Berlin, 1992).
[18]J. Fokkema and P. van den Berg, *Seismic Applications of Acoustic Reciprocity* (Elsevier, Amsterdam, 1993).
[19]E. Kamen, *Introduction to Signals and Systems*, 2nd ed. (McMillan, New York, 1987).
[20]A. Poularikas and S. Seely, *Signals and Systems*, 2nd ed. (PWS-KENT, Boston, 1991).
[21]R. van der Veeken and H. Blok, "A configurational filtering and sampling method in the numerical solution of three-dimensional electromagnetic source problems," J. Electromagn. Waves Appl. **4**, 919–943 (1990).
[22]M. Abramowitz and I. Stegun, *Handbook of Mathematical Functions*, 9th ed. (Dover, New York, 1972).
[23]R. Zemp, J. Tavakkoli, and R. Cobbold, "Modeling of nonlinear ultrasound propagation in tissue from array transducers," J. Acoust. Soc. Am. **113**, 139–152 (2003).
[24]A. Grama, A. Gupta, G. Karypis, and V. Kumar, *Introduction to Parallel Computing*, 2nd ed. (Pearson Education Limited, Harlow, England, 2003).
[25]B. Fornberg, *A Practical Guide to Pseudospectral Methods* (Cambridge University Press, Cambridge, 1998).
[26]M. Frigo and S. Johnson, "The design and implementation of FFTW3," Proc. IEEE **93**, 216–231 (2005).
[27]G. Harris, "Review of transient field theory for a baffled planar transducer," J. Acoust. Soc. Am. **70**, 10–20 (1981).
[28]J. Jensen and N. Svendsen, "Calculation of pressure fields from arbitrarily shaped, apodized and excited ultrasound transducers," IEEE Trans. Ultrason. Ferroelectr. Freq. Control **39**, 262–267 (1992).
[29]The behavior of the Neumann iterative solution in combination with the filtered convolution method has been further investigated for a configuration consisting of an acoustically penetrable, homogeneous sphere in a homogeneous background medium (see Ref. 11). For this case analytical results exist. Successive approximations $\delta p^{(j)}=p^{(j)}-p^{(0)}$ of the scattered field have been determined for several radii of the sphere, several values of the wave speed inside the sphere, and several values of the sampling parameter D_F. For those situations in which the successive approximations clearly tend to a limiting result, the RRMS error of this limiting result with respect to the exact solution has been evaluated. From these comparisons, it may be concluded that the limiting results are indeed an approximation of the exact solution, and become better for higher values of D_F. The latter is explained by the fact that the observed error is almost entirely caused by the spatial filtering of the contrast function of the relatively small sphere.