A Coupled Finite Difference – Gaussian Beam Method for High Frequency Wave Propagation

Nicolay M. Tanushev, Richard Tsai, and Björn Engquist

Abstract Approximations of geometric optics type are commonly used in simulations of high frequency wave propagation. This form of technique fails when there is strong local variation in the wave speed on the scale of the wavelength or smaller. We propose a domain decomposition approach, coupling Gaussian beam methods where the wave speed is smooth with finite difference methods for the wave equations in domains with strong wave speed variation. In contrast to the standard domain decomposition algorithms, our finite difference domains follow the energy of the wave and change in time. A typical application in seismology presents a great simulation challenge involving the presence of irregularly located sharp inclusions on top of a smoothly varying background wave speed. These sharp inclusions are small compared to the domain size. Due to the scattering nature of the problem, these small inclusions will have a significant effect on the wave field. We present examples in two dimensions, but extensions to higher dimensions are straightforward.

1 Introduction

In this paper, we consider the scalar wave equation,

$$\begin{aligned} \Box u = u_{tt} - c^2(x)\triangle u = 0 \qquad & (t,x) \in [0,T] \times \mathbb{R}^d \\ u(0,x) = f(x) & \\ u_t(0,x) = g(x)\ , & \end{aligned} \tag{1}$$

Nicolay M. Tanushev (nicktan@math.utexas.edu),
Richard Tsai (ytsai@math.utexas.edu),
Björn Engquist (engquist@math.utexas.edu)
The University of Texas at Austin, Department of Mathematics,
1 University Station, C1200, Austin, TX 78712.

where d is the number of space dimensions. We will mainly focus on $d = 2$, though the extension of the methods presented here to three or more spatial dimensions is straight forward. The wave equation (1) is well-posed in the energy norm,

$$\|u(t,\cdot)\|_E^2 = \int_{\mathbb{R}^d} \left[\frac{|u_t(t,x)|^2}{c^2(x)} + |\nabla u(t,x)|^2 \right] dx \ , \tag{2}$$

and it is often useful to define the point-wise energy function,

$$E[u](t,x) = \frac{|u_t(t,x)|^2}{c^2(x)} + |\nabla u(t,x)|^2 \ , \tag{3}$$

and the energy inner product,

$$< u, v >_E = \int_{\mathbb{R}^d} \left[\frac{u_t(t,x)\bar{v}_t(t,x)}{c^2(x)} + \nabla u(t,x) \cdot \nabla \bar{v}(t,x) \right] dx$$

High frequency solutions to the wave equation (1) are necessary in many scientific applications. While the equation has no scale, "high frequency" in this case means that there are many wave oscillations in the domain of interest and these oscillations are introduced into the wave field from the initial conditions. In simulations of high frequency wave propagation, direct discretization methods are notoriously computationally costly and typically asymptotic methods such as geometric optics [4], geometrical theory of diffraction [8], and Gaussian beams [2, 5, 6, 7] are used to approximate the wave field. All of these methods rely on the underlying assumption that the wave speed $c(x)$ does not significantly vary on the scale of the wave oscillations. While there are many interesting examples in scientific applications that satisfy this assumption, there are also many cases in which it is violated, for example in seismic exploration, where inclusions in the subsurface composition of the earth can cause the wave speed to vary smoothly on the scale of seismic wavelengths or even smaller scales. In this paper, we are interested in designing coupled simulation methods that are both fast and accurate for domains in which the wave speed is rapidly varying in some subregions of the domain and slowly varying in the rest.

In typical domain decomposition algorithms, the given initial-boundary value problem (IBVP) is solved using numerical solutions of many similar IBVPs on smaller subdomains with fixed dimensions. The union of these smaller domains constitutes the entire simulation domain. In our settings, there are two major differences to the case above. First, the equations and numerical methods in the subdomains are different: we have subdomains in which the wave equation is solved by a finite difference method while in other subdomains the ODEs defined by the Gaussian beam method are solved. Second, we consider situations in which the wave energy concentrates on small subregions of the given domain, so our domain decomposition method requires

subdomains which follow the wave energy propagation and thus change size and location as a function of time. Since our method couples two different models of wave propagation, we will refer to it as the hybrid method. These types of methods are also often called heterogeneous domain decomposition [11]. We will describe how information is exchanged among the subdomains as well as how to change the subdomain size without creating instability and undesired numerical effects.

Our strategy will be to use an asymptotic method in subregions of the domain that satisfy the slowly varying sound speed assumption and a local direct method based on standard centered differences in subregions that do not. This hybrid domain decomposition approach includes three steps. The first is to translate a Gaussian beam representation of the high frequency wave field to data for a full wave equation finite difference simulation. Since a finite difference method needs the values of the solution on two time levels, this coupling can be accomplished by simply evaluating the Gaussian beam solution on the finite difference grid. The next step is to perform the finite difference simulation of the wave equation in an efficient manner. For this, we design a local finite difference method that simulates the wave equation in a localized domain, which moves with the location of a wave energy. Since this is a major issue, we have devote a section of this paper to its description and provide some examples. The last step is to translate a general wave field from a finite difference simulation to a superposition of Gaussian beams. To accomplish this, we use the method described in [14] for decomposing a general high frequency wave field $(u, u_t) = (f, g)$ into a sum of Gaussian beams. The decomposition algorithm is a greedy iterative method. At the $(N + 1)$ decomposition step, a set of initial values for the Gaussian beam ODE system is found such that the Gaussian beam wave field given by these initial values will approximates the residual between the wave field (f, g) and the wave field generated by previous (N) Gaussian beams at a fixed time. These new initial values are directly estimated from the residual wave field and are then locally optimized in the energy norm using the Nelder-Mead method [10]. The procedure is repeated until a desired tolerance or maximum number of beams is reached.

Since Gaussian beam methods are not widely known, we begin with a condensed description of Gaussian beams. After this presentation, we give two examples that show the strengths and weaknesses of using Gaussian beams. We develop the local finite difference method as a stand alone method for wave propagation. Finally, we combine Gaussian beams and the local finite difference method to form the hybrid domain decomposition method. We present two examples to show the strength of the hybrid method.

2 Gaussian beams

Since Gaussian beams play a central role in the hybrid domain decomposition method, we will briefly describe their construction. For a general construction and analysis of Gaussian beams, we refer the reader to [12, 13, 9].

Gaussian beams are approximate high frequency solutions to linear PDEs which are concentrated on a single ray through space–time. They are closely related to geometric optics. In both approaches, the solution of the PDE is assumed to be of the form $a(t,x)e^{ik\phi(t,x)}$, where k is the large high frequency parameter, a is the amplitude of the solution, and ϕ is the phase. Upon substituting this ansatz into the PDE, we find the eikonal and transport equations that the phase and amplitude functions have to satisfy, respectively. In geometric optics ϕ is real valued, while in Gaussian beams ϕ is complex valued. To form a Gaussian beam solution, we first pick a characteristic ray for the eikonal equation and solve a system of ODEs in t along it to find the values of the phase, its first and second order derivatives and amplitude on the ray. To define the phase and amplitude away from this ray to all of space–time, we extend them using a Taylor polynomial. Heuristically speaking, along each ray we propagate information about the phase and amplitude that allows us to reconstruct them locally in a Gaussian envelope.

For the wave equation, the system of ODEs that define a Gaussian beam are

$$\begin{aligned}
\dot{\phi}_0(t) &= 0 \ , \\
\dot{y}(t) &= -c(y(t))p(t)/|p(t)| \ , \\
\dot{p}(t) &= |p(t)|\nabla c(y(t)) \ , \\
\dot{M}(t) &= -A(t) - M(t)B(t) - B^{\mathsf{T}}(t)M(t) - M(t)C(t)M(t) \ , \\
\dot{a_0}(t) &= a_0(t)\left(-\frac{p(t)}{2|p(t)|}\cdot\frac{\partial c}{\partial x}(y(t)) - \frac{p(t)\cdot M(t)p(t)}{2|p(t)|^3} + \frac{c(y(t))\mathrm{Tr}[M(t)]}{2|p(t)|}\right) \ ,
\end{aligned}$$

where

$$\begin{aligned}
A(t) &= -|p(t)|\frac{\partial^2 c}{\partial x^2}(y(t)) \ , \\
B(t) &= -\frac{p(t)}{|p(t)|}\otimes\frac{\partial c}{\partial x}(y(t)) \ , \\
C(t) &= -\frac{c(y(t))}{|p(t)|}\left(Id_{d\times d} - \frac{p(t)\otimes p(t)}{|p(t)|^2}\right) \ .
\end{aligned}$$

The quantities $\phi_0(t)$ and $a_0(t)$ are scalar valued, $y(t)$ and $p(t)$ are in $\mathbb{R}^d$, and $M(t)$, $A(t)$, $B(t)$, and $C(t)$ are $d\times d$ matrices. Given initial values, the solution to this system of ODEs will exists for $t\in[0,T]$, provided that $M(0)$ is symmetric and its imaginary part is positive definite. Furthermore, $M(t)$ will remain symmetric with a positive definite imaginary part for $t\in[0,T]$.

For a proof, we refer the reader to [12]. Under the restriction on $M(0)$, the ODEs allow us to define the phase and amplitude for the Gaussian beam using:

$$\begin{aligned} \phi(t,x) &= \phi_0(t) + p(t)\cdot[x-y(t)] + \frac{1}{2}[x-y(t)]\cdot M(t)[x-y(t)] \\ a(t,x) &= a_0(t)\ . \end{aligned} \tag{4}$$

Furthermore, since $\dot{\phi}_0(t) = 0$, for fixed k, we can absorb this constant phase shift into the amplitude and take $\phi_0(t) = 0$. Thus, the Gaussian beam solution is given by

$$v(t,x) = a(t,x)e^{ik\phi(t,x)}\ . \tag{5}$$

We will assume that the initial values for these ODEs are given and that they satisfy the conditions on $M(0)$. The initial values for the ODEs are tied directly to the Gaussian beam wave field at $t = 0$, $v(0,x)$ and $v_t(0,x)$. As can be easily seen, the initial conditions for the Gaussian beam will not be of the general form of the conditions for the wave equation given in (1). However, using a decomposition method such as the methods described in [14] or [1], we can approximate the general high frequency initial conditions for (1) as a superposition of individual Gaussian beams. Thus, for the duration of this paper, we will assume that the initial conditions for the wave equation (1) are the same as those for a single Gaussian beam:

$$\begin{aligned} u(0,x) &= a(0,x)e^{ik\phi(0,x)} \\ u_t(0,x) &= \left[a_t(0,x) + ik\phi_t(0,x)a(0,x)\right]e^{ik\phi(0,x)}\ . \end{aligned} \tag{6}$$

Note that $a_t(0,x)$ and $\phi_t(0,x)$ are directly determined by the Taylor polynomials (4) and the ODEs above.

3 Motivating Examples

We begin with an example that shows the strengths of using Gaussian beams and, with a small modification, the shortcomings. Suppose that we consider the wave equation (1) in two dimension for $(t, x_1, x_2) \in [0, 2.5] \times [-1.5, 1.5] \times [-3, 0.5]$, sound speed $c(x) = \sqrt{1 - 0.05x_2}$, and the Gaussian beam initial conditions given in (6) with,

$$\begin{aligned} \phi(0,x) &= (x_2 - 1) + i(x_1 - 0.45)^2/2 + i(x_2 - 1)^2/2\ , \\ a(0,x) &= 1\ . \end{aligned}$$

We take the high frequency parameter $k = 100$. To obtain a numerical solution to the wave equation (1), we can use either a direct method or the

Gaussian beam method. As the direct method, we use the standard second order finite difference method based on the centered difference formulas for both space and time:

$$\begin{aligned} &\frac{u^{n+1}_{\ell,m} - 2u^n_{\ell,m} + u^{n-1}_{\ell,m}}{\Delta t^2} \\ &\quad = c^2_{\ell,m} \left[\frac{u^n_{\ell+1,m} - 2u^n_{\ell,m} + u^n_{\ell-1,m}}{\Delta x^2} + \frac{u^n_{\ell,m+1} - 2u^n_{\ell,m} + u^n_{\ell,m-1}}{\Delta y^2} \right] , \end{aligned} \tag{7}$$

where n is the time level index, ℓ and m are the x and y spatial indices respectively.

Since we need to impose artificial boundaries for the numerical simulation domain, we use first order absorbing boundary conditions (ABC) [3]. The first order ABC amount to using the appropriate one-way wave equation,

$$u_t = \pm c(x,y) u_x \text{ or } u_t = \pm c(x,y) u_y \ , \tag{8}$$

on each of the boundaries, so that waves are propagated out of the simulation domain and not into it. For example, on the left boundary, $x = -1.5$, we use $u_t = c u_x$ with upwind discretization,

$$\frac{u^{n+1}_{\ell,m} - u^n_{\ell,m}}{\Delta t} = c_{\ell,m} \left[\frac{u^n_{\ell+1,m} - u^n_{\ell,m}}{\Delta x} \right] , \tag{9}$$

for ℓ equal to its lowest value.

To resolve the oscillations, using 10 points per wavelength, for this particular domain size and value for k, we need roughly 500 points in both the x_1 and x_2 directions. However, to maintain low numerical dispersion for the finite difference solution, we need to use a finer the grid. The grid refinement will the given in terms of the coarse, 10 points per wavelength, grid. For example, a grid with a refinement factor of 3 will have 30 points per wavelength. Note that such grid refinement is not necessary for the Gaussian beam solution. Thus, while we compute the finite difference solution on the refined grid, we only use the refined solution values on the coarser grid for comparisons. For determining the errors in each solution, we compare with the "exact" solution computed using the finite difference method with a high refinement factor of 10.

For this particular example, the sound speed, the finite difference solution and Gaussian beam solution at the final time are shown in Figure 1. In order to have a meaningful comparison, the grid refinement for the finite difference solution was chosen so that the errors in the finite difference solution are comparable to the ones in the Gaussian beam solution. Both the accuracy and computation times are shown in Table 1. The Gaussian beam solution was computed more than 3500 times faster than the finite difference solution and the total error for both the Gaussian beam and the finite difference solution is $\approx 10\%$. Near the center of the beam, where the Gaussian beam envelope is

greater than 0.25, the Gaussian beam solution is slightly more accurate with a local error of $\approx 7\%$. The Gaussian beam solution is an asymptotic solution, thus its error decreases for larger values of k. In terms of complexity analysis, as we are using a fixed number of points per wavelength to represent the wave field, the Gaussian beam solution is computed in $\mathcal{O}(1)$ steps and evaluated on the grid in $\mathcal{O}(k^2)$. The finite difference solution is computed in $\mathcal{O}(k^3)$ steps. Additionally, for larger values of k, we would need to increase the grid refinement for the finite difference solution in order to maintain the same level of accuracy as in the Gaussian beam solution. Therefore, it is clear why the Gaussian beam solution method is advantageous for high frequency wave propagation.

	t=0.625	t=1.25	t=1.875	t=2.5	Loc Err	C Time
FD	1.9%	3.8%	5.6%	7.3%	7.4%	7773.1
GB	2.4%	4.8%	7.2%	9.7%	7.0%	1.6

Table 1 Comparisons of the finite difference (FD) method and Gaussian beam (GB) method with sound speed with no inclusion. Shown are the total error for each method in the energy norm as a percent of the total energy at each time, the local error (Loc Err) as a percent of the local energy at $t = 2.5$, and the total computational time (C Time) for obtaining the solution at each time. The local error is computed near the beam center, where the Gaussian envelope is greater than 0.25. The finite difference solution is computed with a refinement factor of 6.

Now, suppose that we modify the sound speed to have an inclusion, so that the sound speed changes on the same scale as the wave oscillations as shown in Figure 1 and that we use the same initial conditions as before. The inclusion is positioned in such a way, so that the ray mostly avoids the inclusion, while the wave field on the left side of the ray interacts with the inclusion. Since all of the quantities that define the Gaussian beam are computed on the ray, the Gaussian beam coefficients are similar to the coefficients in the example without the inclusion. However, as can be seen from the full finite difference calculation in Figure 1, the wave field at $t = 2.5$ is very different from the wave field at $t = 2.5$ for the sound speed with no inclusion shown in the same figure. The solution errors shown in Table 2 demonstrate that, while the Gaussian beam computation time is again more than 3500 times faster, the error renders the solution essentially useless. Thus, the Gaussian beam solution is not a good approximation of the exact solution in this case. This, of course, is due to the fact that the asymptotic assumption, that the sound speed is slowly varying, is violated. Therefore, for a sound speed with an inclusion of this form, the Gaussian beam method cannot be used and we have to compute the wave field using a method that does not rely on this asymptotic assumption.

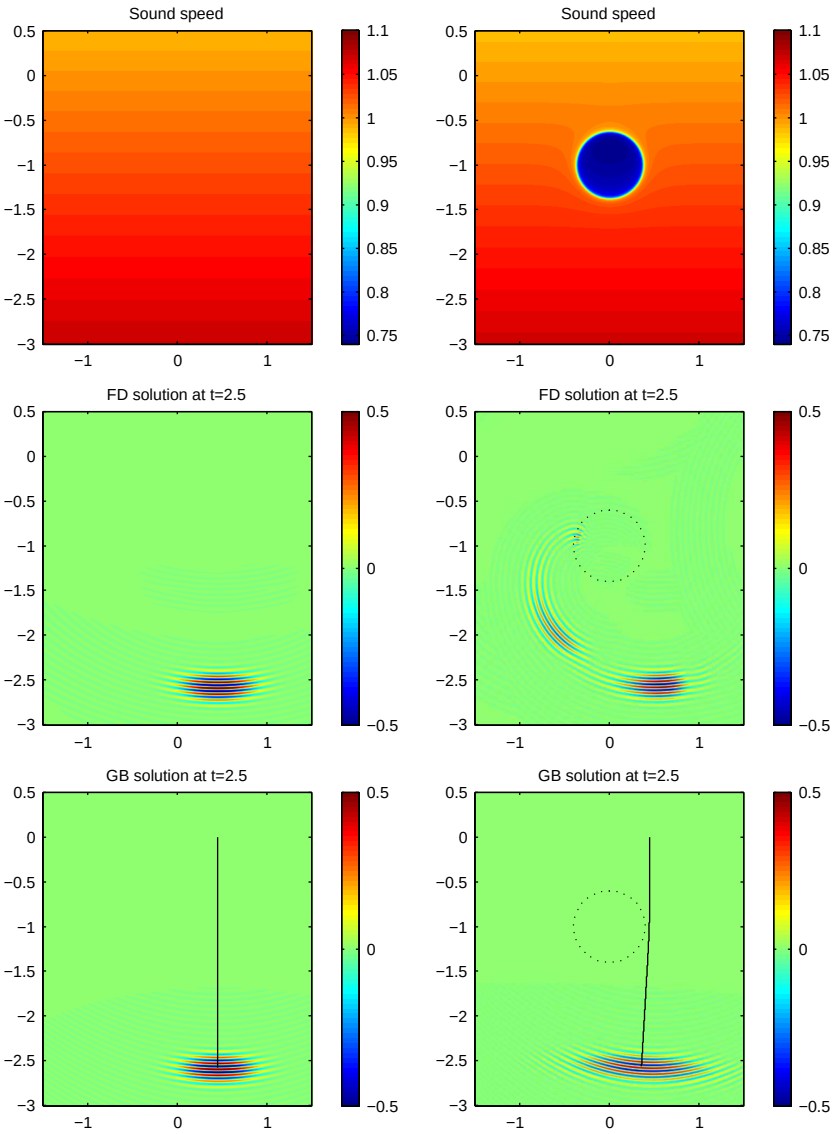

Fig. 1 The first column shows the wave field for simulations with sound speed without an inclusion: sound speed, the finite difference (FD) solution at the final time, and the Gaussian beam (GB) solution at the final time. The second column shows the same graphs for simulations with sound speed containing an inclusion. The line shows the ray for the Gaussian beam. At $t = 0$, the Gaussian beam is centered at the beginning of the line and at $t = 2.5$, it is centered at the end of the line. The dotted circle outlines the location of the inclusion in the sound speed. For each of the wave fields, only the real part is shown.

	t=0.625	t=1.25	t=1.875	t=2.5	Loc Err	C Time
FD	1.9%	3.9%	5.6%	7.0%	7.3%	7717.8
GB	6.1%	94.5%	91.2%	90.9%	43.9%	1.5

Table 2 Comparisons of the finite difference (FD) method and Gaussian beam (GB) method for a sound speed with inclusion. Shown are the total error for each method in the energy norm as a percent of the total energy at each time, the local error (Loc Err) as a percent of the local energy at $t = 2.5$, and the total computational time (C Time) for obtaining the solution at each time. The local error is computed near the beam center, where the Gaussian envelope is greater than 0.25. The finite difference solution is computed with a refinement factor of 6.

4 Local Finite Difference Method

By examining the example in the previous section, it is clear that a large portion of the computational time for the finite difference solution is spent simulating the wave equation where the solution is nearly zero. To exploit this property of the solution, we propose to use finite differences to compute the solution only locally where the wave energy is concentrated. Since the wave energy propagates in the domain, the region in which we carry out the local wave equation simulation must also move with the waves. We emphasize that we are not using Gaussian beams at this stage.

To be more precise, we propose to simulate the wave equation in a domain $\Omega(t)$, that is a function of time and at every t, $\Omega(t)$ contains most of the wave energy. For computational ease, we select $\Omega(t)$ to be a rectangular region.The initial simulation domain $\Omega(0)$ is selected from the initial data by thresholding the energy function (3) to contain most of the wave energy. Since solutions of the wave equation (1) have finite speed of propagation, the energy moves at the speed of wave propagation and thus the boundaries of $\Omega(t)$ do not move too rapidly. In terms of finite difference methods, this means that if we ensure that the Courant-Friedrichs-Lewy (CFL) condition is met, the boundaries of $\Omega(t)$ will not move by more than a spatial grid point between discrete time levels t and $t+\Delta t$. Whether $\Omega(t)$ increases or decreases by one grid point (or stays the same) at time level $t + \Delta t$ is determined by thresholding the energy function (3) of u at time level t near the boundary of $\Omega(t)$.

Using the standard second order finite difference method, we discretize the wave equation (1) using a centered in time, centered in space finite difference approximation (7). Since the solution is small near the boundary of $\Omega(t)$, there are several different boundary conditions that we could implement to obtain a solution. The easiest and most straightforward approach is to simply use Dirichlet boundary conditions with $u = 0$. Another approach is to use absorbing boundary conditions. We investigate the case where absorbing boundary conditions are applied to a single layer of grid nodes immediately neighboring the outer most grid nodes of $\Omega(t)$ (single layer ABC) and absorbing boundary conditions are applied again to the layer of grid nodes

immediately neighboring the first ABC layer (double layer ABC). For example, for the depicted grid nodes in Figure 2, $u^{n+1}_{\ell+1,m}$ and $u^{n+1}_{\ell,m}$ are computed by

$$u^{n+1}_{\ell+1,m} = u^n_{\ell+1,m} + c_{\ell+1,m}\frac{\Delta t}{\Delta x}\left[u^n_{\ell+2,m} - u^n_{\ell+1,m}\right]$$
$$u^{n+1}_{\ell,m} = u^n_{\ell,m} + c_{\ell,m}\frac{\Delta t}{\Delta x}\left[u^n_{\ell+1,m} - u^n_{\ell,m}\right] \ .$$

For both Dirichlet and absorbing boundary conditions, when the domain $\Omega(t)$ is expanding, the finite difference stencils will need to use grid nodes that are outside of $\Omega(t)$ and the boundary layers. We artificially set the wave field to be equal to zero at such grid nodes and we will refer to them as "reclaimed grid nodes". Figure 2 shows the domain of influence of the reclaimed nodes for the Dirichlet boundary conditions and the double layer ABC. In this figure, solid lines connect the reclaimed nodes with nodes whose values are computed directly using the reclaimed nodes. Dashed lines connect the reclaimed nodes with nodes whose values are computed using the reclaimed nodes, but through the values of another node. Finally, dotted lines indicate one more level in the effect of the reclaimed nodes. The point of using double layer ABC is to minimize the influence of the reclaimed nodes, as can be seen in Figure 2. Note that there are no solid line connections between the reclaimed nodes and the nodes in $\Omega(t)$ for double layer ABC. Furthermore, the artificial Dirichlet boundary conditions reflect energy back into the computational domain $\Omega(t)$ which may make it larger compared to $\Omega(t)$ for the solution obtained by double layer ABC as shown in Figure 3.

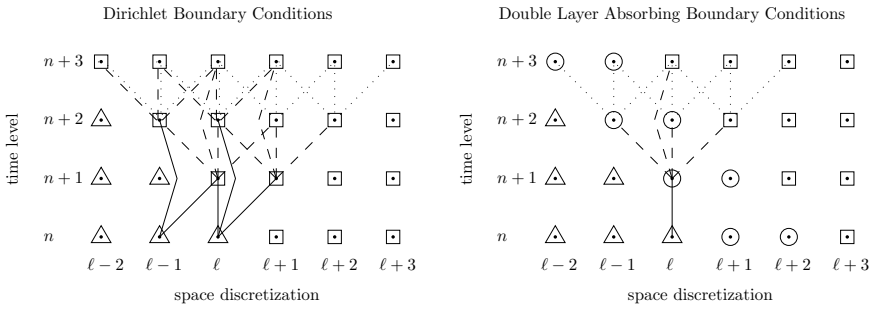

Fig. 2 A comparison between the domains of influence of the reclaimed grid nodes for Dirichlet and double layer absorbing boundary conditions. The wave field is computed at the square grid nodes using centered in time centered in space finite differences and at the circle grid nodes using absorbing boundary conditions. The triangle grid nodes are the reclaimed grid nodes with artificial zero wave field. The lines indicate how the finite differences propagate these artificially values from the n-th time level to later time levels.

Finally, we note that due to finite speed of wave propagation, we can design boundary conditions that will not need reclaimed grid nodes. However, these

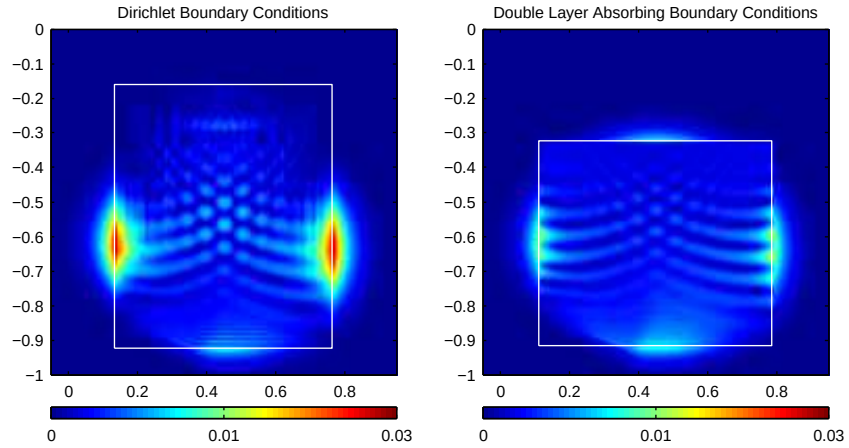

Fig. 3 A comparison between Dirichlet boundary conditions and double layer absorbing boundary conditions for the local finite difference method. The absolute value of the difference between each solution and the finite difference solution for the full domain is plotted at time $t = 0.625$. The domain $\Omega(0.625)$ is outlined in white. Note that overall the double layer absorbing boundary conditions solution is more accurate than the Dirichlet boundary condition solution. Also, note that $\Omega(0.625)$ is smaller for the double layer absorbing boundary conditions.

boundary conditions may have a finite difference stencil that spans many time levels and this stencil may need to change depending on how $\Omega(t)$ changes in time. Numerically, we observed a large improvement when using double layer ABC instead of Dirichlet boundary conditions. However, using triple or quadruple layer ABC did not give a significant improvement over the double layer ABC. Thus, for computational simplicity, we use the above double layer absorbing boundary conditions for the simulations that follow.

Using the local finite difference method, we compute the solution to the wave equation (1) as in the previous section for the example with a sound speed with inclusion, using a refinement factor of 6. To determine $\Omega(0)$, we threshold the energy function (3) at $1/100$ of its maximum. For computational time comparison, we also compute the full finite difference solution, also with a refinement factor of 6. These parameters were chosen so that the final error is $\approx 7\%$ and comparable for both solutions. The wave field, along with $\Omega(t)$, are shown in Figure 4 at $t = \{0, 0.625, 1.25, 1.875, 2.5\}$. The comparisons of accuracy and computation time between the local and full finite difference solutions are shown in Table 3. The error in both solutions is equivalent, but the local finite difference solution in computed 5 times faster. Furthermore, if the local finite difference method is used to simulate the wave field from a Gaussian beam, we need $\mathcal{O}(k)$ steps in time as in the full finite difference method, but the local finite difference method requires $\mathcal{O}(k)$ grid points in space as opposed to $\mathcal{O}(k^2)$ grid points that the full finite difference method

requires. This is because the energy from a Gaussian beam is concentrated in a $k^{-1/2}$ neighborhood of its center and this is a two dimensional example.

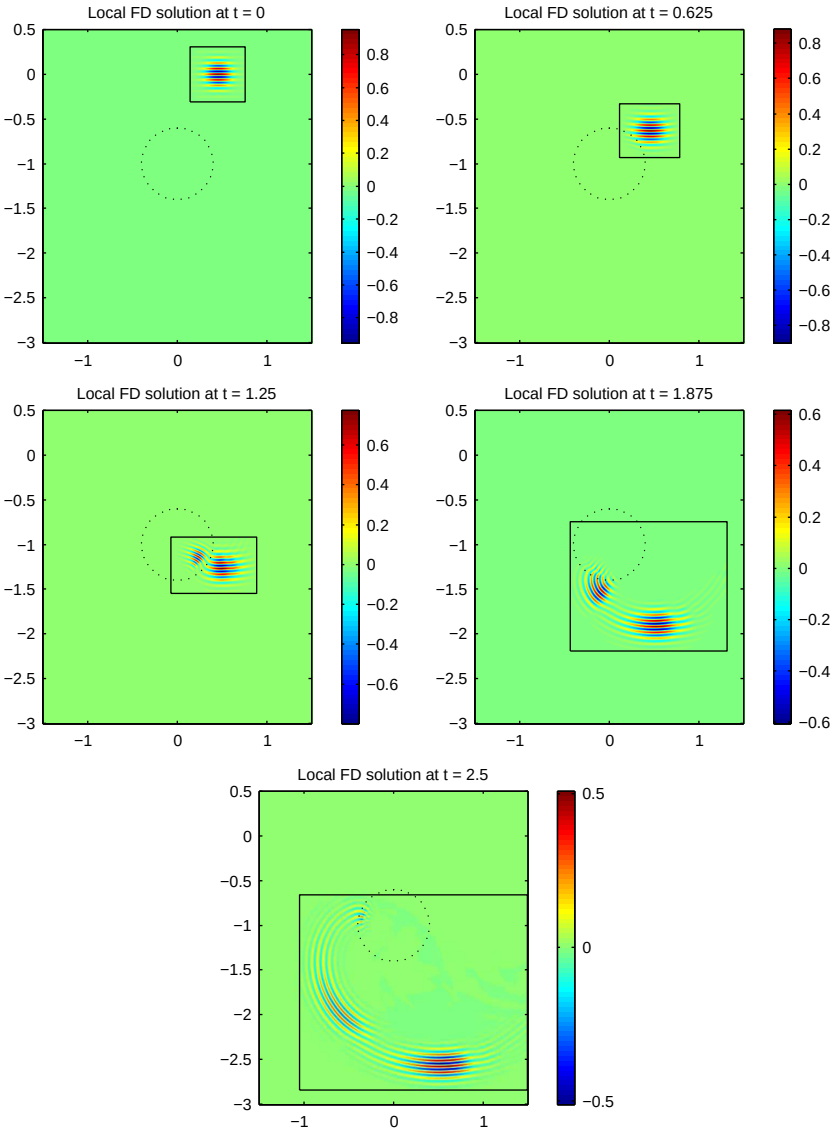

Fig. 4 This figure shows the wave field computed using the local finite difference method for the sound speed with inclusion. The black rectangle outlines the local computational domain, $\Omega(t)$, and the dotted circle outlines the location of the inclusion in the sound speed. Only the real part of the wave fields is shown.

	t=0.625	t=1.25	t=1.875	t=2.5	C Time
FD	1.9%	3.9%	5.6%	7.0%	7717.8
LFD	2.4%	4.4%	6.0%	7.3%	1535.5

Table 3 Comparisons of the full finite difference (FD) method and the local finite difference (LFD) method with sound speed with inclusion. Shown are the total error for each method in the energy norm in as a percent of the total energy and the total computational time (C Time) for obtaining the solution at $t = \{0.625, 1.25, 1.875, 2.5\}$.

Finally, we remark that if instead of finding one rectangle that contains the bulk of the energy we found several, the solution in each of these rectangles can be computed independently. On a parallel computer, this would give another advantage over full finite difference simulations, as there is no need for information exchange between the computations on each rectangle, even if these rectangles overlap. The linear nature of the wave equation allows for the global solution to be obtained by simply adding the solutions from each of the separate local finite difference simulations. Furthermore, the generalization to more than two dimensions is straight forward and the computational gain is even greater in higher dimensions.

5 Hybrid Method

Upon further examination of the inclusion example in Section 3 and the wave field simulations in Section 4, we note that the Gaussian beam solution has small error for some time initially (see Table 2) and that after the wave energy has interacted with the inclusion in the sound speed, it again appears to have Gaussian beam like characteristics (see Figure 4, $t > 2$). We can immediately see the effect of the large variation of the sound speed on the wave field. The large gradient roughly splits the wave field into two components, one that continues on nearly the same path as before and one that is redirected to the side. This also shows why the Gaussian beam solution is not a very good approximation. For a single Gaussian beam to represent a wave field accurately, the wave field has to stay coherent; it cannot split into two or more separate components. However, once the wave field has been split into several components by the inclusion, it will propagate coherently until it reaches another region of large sound speed variation. By following the propagation of wave energy in time, while it is near a region of high sound speed variation, we employ the local finite different method and the Gaussian beam method otherwise.

To be able to use such a hybrid method, we need to be able to couple the two different simulation methods. Switching from a Gaussian beam description to a local finite difference description is straightforward. The local finite difference requires the wave field at a time t and $t + \Delta t$, which can be obtained simply by evaluating the Gaussian beam solution on the finite

difference grid. The opposite, moving from a local finite difference to a Gaussian beam description, is more difficult to accomplish. For this step we use the decomposition algorithm given in [14]. As discussed in the introduction, this decomposition method is a greedy iterative method. At each iteration the parameters for a single Gaussian beam are estimated and then locally optimized using the Nelder-Mead algorithm [10]. The method is then iterated over the residual wave field. The decomposition is complete when a certain tolerance is met or a maximum number of Gaussian beams is reached. For completeness, we give the algorithm of [14] below.

1. With $n = 1$, let (u^n, u_t^n) be the wave field at a fixed t and suppress t to simplify the notation.
2. Find a candidate Gaussian beam:
   - Estimate Gaussian beam center:
     $\rightarrow$ Let $\tilde{y}^n = \arg\max\{E[u^n](y)\}$ (see equation (3)).
   - Estimate propagation direction:
     $\rightarrow$ Let $G(x) = \exp(-k|x - \tilde{y}^n|^2/2)$
     $\rightarrow$ Let $p^n = \arg\max\{|\mathcal{F}[u^n(x)G(x)]| + |\mathcal{F}[u_t^n(x)G(x)/k]|\}$, with $\mathcal{F}$ the scaled Fourier transform, $\{x \rightarrow kp\}$
     $\rightarrow$ Let $\tilde{\phi}_t^n = c(y^n)|\tilde{p}^n|$
   - Let $\tilde{M}^n = iI$, with I the identity matrix.
3. Minimize the difference between the Gaussian beam and u^n in the energy norm using the Nelder–Mead method with $(\tilde{y}^n, \tilde{\phi}_t^n, \tilde{p}^n, \tilde{M}^n)$ as the initial Gaussian beam parameters:
   - Subject to the constraints, $\mathrm{Im}\{M\}$ is positive definite, entries of M are less than $\sqrt{k}$ in magnitude, $1/\sqrt{k} \leq |p| \leq \sqrt{k}$, and $|\phi_t|^2 = c^2(y)|p|^2$, let
   $$(y^n, \phi_t^n, p^n, M^n) = \arg\min \left\{ \left|\left| u^n - \frac{< u^n, B >_E}{||B||_E^2} B \right|\right|_E^2 \right\} ,$$
   where B be the Gaussian beam defined by the initial parameters $(y^n, \phi_t^n, p^n, M^n)$ and amplitude 1 (see equations (4) and (5)).
   - Let $B^n(x,t)$ be the Gaussian beam defined by the initial parameters $(y^n, \phi_t^n, p^n, M^n)$ and amplitude 1.
   - Let $a^n = \frac{<u^n, B^n>_E}{||B^n||_E^2}$.
4. The n-th Gaussian beam is given by the parameters $(y^n, \phi_t^n, p^n, M^n, a^n)$. Subtract its wave field:
   $$u^{n+1} = u^n - a^n B^n \text{ and } u_t^{n+1} = u_t^n - a^n B_t^n \ .$$
5. Re-adjust the previous $n - 1$ beams:
   - For the j-th beam, let $w = u^{n+1} + a^j B^j$ and repeat step 3 with $u^n = w$, $n = j$, and $(y^j, \phi_t^j, p^j, M^j)$ as the initial Gaussian beam parameters.

- Let $u^{n+1} = w - a^j B^j$

6. Re-adjust all beam amplitudes together
   - Let Λ be the matrix of inner products $\Lambda_{j\ell} =< B^\ell, B^j >_E$, and $b^j =< u^1, B^j >_E$
   - Solve $\Lambda a = b$ and let $u^{n+1} = u^1 - \sum_{j=1}^n a^j B^j$
7. Repeat steps starting with step 2, until $||u^{n+1}||_E$ is small or until a prescribed number of Gaussian beams is reached.

The final step in designing the hybrid method is deciding when and where to use which method. By looking at the magnitude of the gradient of the sound speed and the value of k, we can decompose the simulation domain into two subdomains D_G and D_L, which represent the Gaussian beam, small sound speed gradient, subdomain and the local finite difference, large gradient, subdomain respectively. When the Gaussian beam ray enters D_L, we switch from the Gaussian beam method to the local finite difference method. Deciding when to switch back to a Gaussian beam description is again more complicated. One way is to monitor the energy function (3) and when a substantial portion of it is supported in D_G, we use the decomposition method to convert that part of the energy into a superposition of a few Gaussian beams. Since calculating the energy function is computationally expensive, it should not be done at every time level of the local finite difference simulation. From the sound speed and size of D_L, we can estimate a maximum speed of propagation for the wave energy, thus a minimum time to exit D_L, and use that as a guide for evaluating the energy function. Additionally, we can look at the overlap between D_G and the local finite difference simulation domain $\Omega(t)$ as a guide for checking the energy function. A more crude, but faster, approach is to use the original ray to estimate the time that it takes for the wave energy to pass through D_L. We use this approach in the examples below. Furthermore, we note that the linearity property of the wave equation allows us to have a joint Gaussian beam and local finite difference description of the wave field. We can take the part of the local finite difference wave field in D_G and represent it as Gaussian beams. If there is a significant amount of energy left in D_L, we propagate the two wave fields concurrently one using Gaussian beams and the other using the local finite difference method. The total wave field is then the sum of the Gaussian beam and local finite difference wave fields.

There are two advantages of the hybrid method over the full and local finite difference methods. One is a decrease in simulation time. The other is due to the particular application to seismic exploration. For seismic wave fields, the ray based nature of Gaussian beams provides a connection between the energy on the initial surface and its location at the final time. Furthermore, this energy is supported in a tube in space–time and thus it only interacts with the sound speed inside this tube. Unfortunately, for finite difference based methods there is only the domain of dependence and this set can be

quite large compared to the Gaussian beam space–time tube. For example, if the sound speed model is modified locally, only Gaussian beams that have space–time tubes that pass through the local sound speed modifications will need to be re-computed to obtain the total wave field. In contrast, a local sound speed modification requires that the entire finite difference solution be re-computed. For the hybrid method, if we decompose the wave field in single beam whenever we switch back to the Gaussian beam description then at any given time, we will either have a Gaussian beam wave field or a local finite difference wave field. After the simulation is complete we can interpolate the Gaussian beam coefficients to times for which the wave field is given by the local finite difference. Note that the resulting interpolated wave field will not satisfy the wave equation, however we will once again have a space–time tube that follows the energy propagation. Thus, we are interested in using the hybrid method to obtain a one beam solution that approximates the wave field better than the Gaussian beam method.

5.1 Example: Double Slit Experiment

In the simplest version of the Hybrid method, we consider an example in which we first use the local finite difference method to solve the wave equation for a given amount of time, then we switch to a Gaussian beam representation of the field. For this example we are interested in simulating the wave field in a double slit experiment, where coherent waves pass through two slits that are spaced closely together and their width is $\mathcal{O}(k^{-1})$, with $k = 50$. In the finite difference method, the slits are implemented as Dirichlet boundary conditions. It is clear that due to the diffraction phenomenon near the two slits, the Gaussian beam method alone will not give an accurate representation of the wave field. The wave field simulated using the hybrid method is shown in Figure 5 and the error and computational time are shown in Table 4. Note that with 14 Gaussian beams, the computational time for the hybrid solution is still a factor of 3 faster than the full finite difference solution and a factor of 2 faster than the local finite difference solution.

	t=1.25	t=2.5	t=3.75	t=5	C Time
FD	5.91%	10.6%	14.8%	19.1%	470
LFD	6.13%	11%	15.8%	19.7%	270
H	6.13%	12.7%	24.2%	33.9%	150

Table 4 Comparisons of the full finite difference (FD), the local finite difference (LFD) and the hybrid (H) methods for the double slit experiment. Shown for each method are the total error in the energy norm in terms of percent of total energy and the total computational time (C Time) for obtaining the solution at $t = \{1.25, 2.5, 3.75, 5\}$. The norms are computed only on $y < 0$, since we are only interested in the wave field that propagates through the two slits.

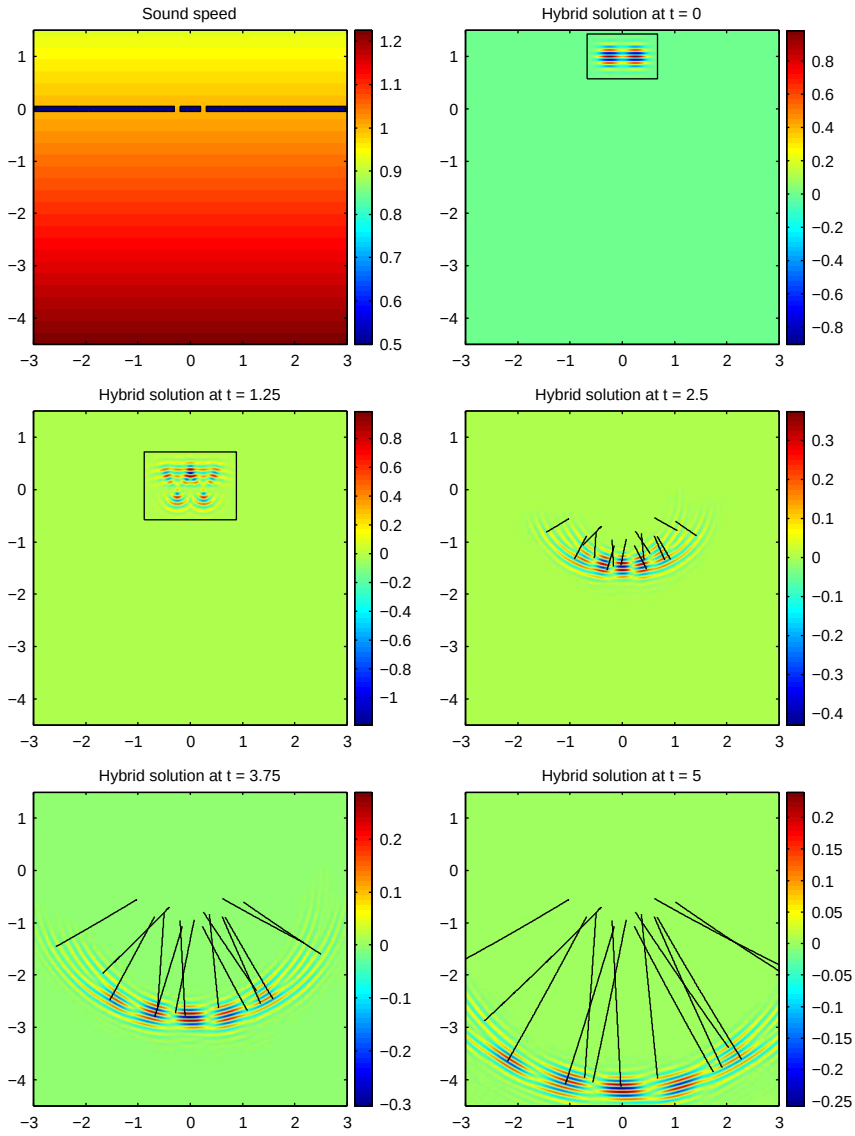

Fig. 5 The wave field obtained by the hybrid method for the double slit experiment. The first panel shows the sound speed and the double slit Dirichlet boundary condition region. The local finite difference domain is outlined by the black rectangle at $t = \{0, 1.25\}$. At $t = \{2.5, 3.75, 5\}$, the black lines indicate the ray for each of the Gaussian beams.

5.2 Example: Sound Speed with Inclusion

Finally, to demonstrate the hybrid method, we apply it to computing the wave field for the sound speed with inclusion and compare it to the previously

discussed methods. For these experiments $k = 100$. The wave field is first computed using Gaussian beams until the beam is close to the inclusion at $t = 0.5$. Then, the solution is propagated with the local finite difference method until most of the wave energy has moved past the inclusion at $t = 2$. The resulting field is then decomposed into one beam (the hybrid one-beam solution) or into two beams (the hybrid two-beam solution) using the decomposition algorithm of [14]. The wave fields for the one and two beam hybrid solutions are shown in Figure 6. The errors and computation times for the methods discussed in this paper are shown in Table 5. The local finite difference calculations are done with a refinement factor of 5 and $\Omega(t)$ is obtained by thresholding the energy function at 1/10 of its maximum. This thresholding was chosen so that the final errors in the local finite difference solution are similar to the error in the hybrid solution making the comparison of the computation times meaningful. The errors for the one and two beam hybrid solutions are $\approx 62\%$ and $\approx 37\%$ respectively at $t = 2.5$. This may seem rather large, but we note that this is a large improvement over the Gaussian beam solution which has an error of $\approx 91\%$. Furthermore, this is a single Gaussian beam approximation of the wave field locally and this wave field is not necessarily of Gaussian beam form. Locally, near the beam centers, the H1 and H2 solutions are more accurate. The computational time for the H1 and H2 hybrid solutions is 2 times faster compared to the local finite difference solution and 10 times faster than the full finite difference solution.

	t=0.675	t=1.25	t=1.875	t=2.5	Loc Err 1	Loc Err 2	C Time
FD	3.3%	6.6%	9.4%	11.8%	12.3%	10.8%	4446.1
GB	6.1%	94.5%	91.2%	90.9%	42.2%	99.9%	1.5
LFD	6.6%	9.6%	11.9%	14.4%	12.4%	10.8%	781.0
H1	3.9%	7.4%	10.2%	62.0%	12.7%	100.0%	401.5
H2	3.9%	7.4%	10.2%	36.7%	12.7%	25.9%	417.9

Table 5 Comparisons of the methods for a sound speed with inclusion. Shown for each method are the total error in the energy norm in terms of percent of total energy at each time, the local errors as a percent of the local energy near the beam center for the first beam (Loc Err 1) and near the second beam center (Loc Err 2), and the total computational time (C Time) for obtaining the solution at each time. The local error is computed near the beam center, where the Gaussian envelope is greater than 0.25. Legend: GB – Gaussian beam, LFD – Local finite difference, H1 – Hybrid method with one beam, H2 – hybrid method with two beams.

6 Conclusion

In this paper, we develop a new hybrid method for high frequency wave propagation. We couple a Gaussian beam approximation of high frequency wave propagation to a local finite difference method in parts of the domains that

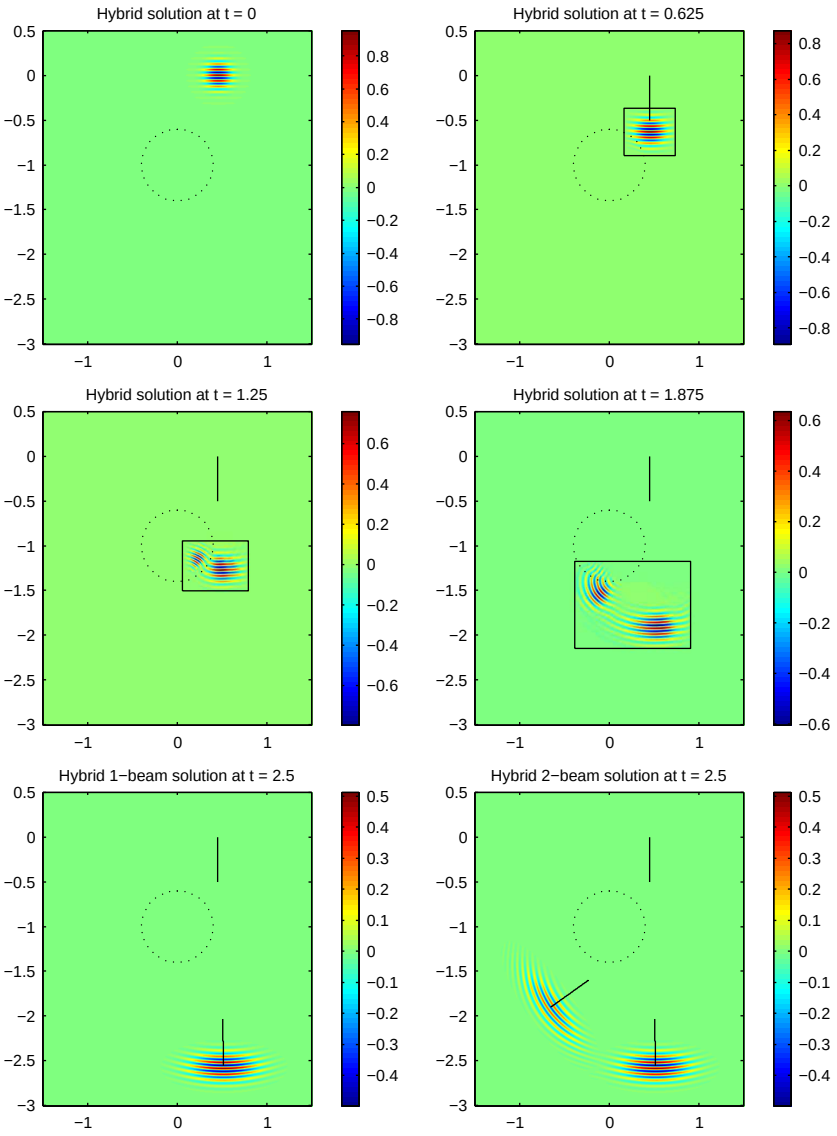

Fig. 6 The wave field for the hybrid H1 and H2 solution. The top two rows show the real part of the wave field which is the same for both the 1–beam and 2–beam hybrid solutions at $t = \{0, 0.625, 1.25, 1.875\}$. Times $t = \{.625, 1.25, 1.875\}$ are during the local finite difference calculation and the black rectangle outline the local finite difference domain $\Omega(t)$. The real part of the wave field for the 1–beam and 2-beam hybrid solutions are shown in the last row at $t = 2.5$. In each panel, the black lines indicate the ray for each of the Gaussian beams.

contain strong variations in the wave speed. The coupling is accomplished either by translating the Gaussian beam representation into a wave field representation on a finite difference grid or by approximating the finite difference solution with a superposition of Gaussian beams. The local finite difference computations are performed on a moving computational domain with absorbing boundary conditions. This direct method is only used at times when a significant portion of the wave field energy is traveling through parts of the domain that contain large variations in the wave speed. The rest of the high frequency wave propagation is accomplished by the Gaussian beam method.

Two numerical test examples show that the hybrid technique can retain the overall computational efficiency of the Gaussian beam method. At the same time the accuracy of the Gaussian beam methods in domains with smooth wave speed field is kept and the accuracy of the finite difference method in domains with strong variation in the wave speed is achieved. Furthermore, the hybrid method maintains the ability to follow the wave energy as it propagates from the initial surface through the domain as in traditional Gaussian beam and other ray based methods.

Acknowledgement

The authors would like to thank Sergey Fomel, Ross Hill and Olof Runborg for helpful discussions and acknowledge the financial support of the NSF. The authors were partially supported under NSF grant No. DMS-0714612. NT was also supported under NSF grant No. DMS-0636586 (UT Austin RTG).

References

1. G. Ariel, B. Engquist, N. Tanushev, and R. Tsai. Gaussian beam decomposition of high frequency wave fields using expectation-maximization. *J. Comput. Phys.*, (to appear), 2011.
2. V. Červený, M. Popov, and I. Pšenčík. Computation of wave fields in inhomogeneous media - Gaussian beam approach. *Geophys. J. R. Astr. Soc.*, 70:109–128, 1982.
3. B. Engquist and A. Majda. Absorbing boundary conditions for the numerical simulation of waves. *Mathematics of Computation*, 31(139):629–651, 1977.
4. B. Engquist and O. Runborg. Computational high frequency wave propagation. *Acta Numer.*, 12:181–266, 2003.
5. S. Gray, Y. Xie, C. Notfors, T. Zhu, D. Wang, and C. Ting. Taking apart beam migration. *The Leading Edge*, Special Section:1098–1108, 2009.
6. R. Hill. Gaussian beam migration. *Geophysics*, 55:1416–1428, 1990.
7. R. Hill. Prestack Gaussian-beam depth migration. *Geophysics*, 66(4):1240–1250, 2001.
8. J. Keller. Geometrical theory of diffraction. *Journal of Optical Society of America*, 52:116–130, 1962.
9. H. Liu and J. Ralston. Recovery of high frequency wave fields for the acoustic wave equation. *Multiscale Modeling & Simulation*, 8(2):428–444, 2009.

10. J. Nelder and R. Mead. A simplex method for function minimization. *The Computer Journal*, 7(4):308–313, 1965.
11. A. Quarteroni, F. Pasquarelli, and A. Valli. Heterogeneous domain decomposition: principles, algorithms, applications. In *Fifth International Symposium on Domain Decomposition Methods for Partial Differential Equations (Norfolk, VA, 1991)*, pages 129–150. SIAM, Philadelphia, PA, 1991.
12. J. Ralston. Gaussian beams and the propagation of singularities. In *Studies in partial differential equations*, volume 23 of *MAA Stud. Math.*, pages 206–248. Math. Assoc. America, Washington, DC, 1982.
13. N. Tanushev. Superpositions and higher order Gaussian beams. *Commun. Math. Sci.*, 6(2):449–475, 2008.
14. N. Tanushev, B. Engquist, and R. Tsai. Gaussian beam decomposition of high frequency wave fields. *J. Comput. Phys.*, 228(23):8856–8871, 2009.