

Geometric Optics in a Phase-Space-Based Level Set and Eulerian Framework¹

Stanley Osher,* Li-Tien Cheng,† Myungjoo Kang,* Hyeseon Shim,* and Yen-Hsi Tsai‡

**Level Set Systems Inc., 1058 Embury Street, Pacific Palisades, California 90272;* †*Department of Mathematics, University of California San Diego, La Jolla, California 92093;* and ‡*Department of Mathematics, University of California Los Angeles, Los Angeles, California 90095*
E-mail: lcheng@math.ucsd.edu

Received October 2, 2001; revised April 17, 2002

We introduce a level set approach for ray tracing and the construction of wavefronts in geometric optics. This is important in a wide variety of applications in wave propagation. Our approach automatically handles the multivalued solutions that appear and automatically resolves the wavefronts. This is achieved through solving for the bicharacteristic strips, whose projection to spatial space gives the wavefronts, in a reduced phase space under a Eulerian and partial differential-equation-based framework. The bicharacteristic strips are represented using a level set approach for handling higher codimensional objects and the partial differential equations responsible for the evolution are reduced forms of the Liouville equations. Results for the two-dimensional case for constant and variable indices of refraction are shown and compared to those of other current methods in the field. Results are also introduced to show the ability to handle reflection and to extend the method to the three-dimensional case. © 2002 Elsevier Science (USA)

1. INTRODUCTION

Geometric optics is an important area of study, especially as an approximation for high-frequency wave propagation. Ray tracing in this setting can be used, for example, to calculate travel times and amplitudes of seismic waves for purposes such as migration and tomography. In fact, calculation of underground seismic waves are of special interest to the oil industry. A number of other areas, such as combustion, crystal growth, and mesh generation, are also influenced by ray-tracing techniques in geometric optics.

Ray tracing can be formulated, in the Lagrangian view, as a solution to a system of ordinary differential equations, the Hamiltonian system, to track the wavefront. This is the

¹ Research supported by AFOSR Grant F49620-01-1-0189.




FIG. 1. Initial set of points forming a circle spread apart during ray tracing.

set of points, forming a curve in \mathbf{R}^2 and a surface in \mathbf{R}^3 , of constant travel time away from the source. The Hamiltonian system gives the trajectories of the rays while transporting other quantities, such as the density and amplitude, along these paths. Thus the standard Lagrangian numerical approach to ray tracing involves representing the wavefront by a set of discrete points and evolving these points along those trajectories. This approach, though easy to understand and to implement, may encounter difficulties in obtaining adequate spatial resolution of the wavefront. This is due to the fact that points close together initially may diverge at later times (see Fig. 1), thus leaving possibly large gaps in the wavefront and adversely affecting resolution. Interpolation steps must be added to keep the wavefront well resolved (see, e.g., [41–43]).

In view of this point, Eulerian approaches are preferable, as they do not follow the points of the wavefront but rather record information about the wavefront on regular and uniform grids in space. Thus spatial resolution of the wavefront is simply controlled by the step size of the underlying grid. However, Eulerian approaches, which instead solve a Hamilton–Jacobi partial differential equation, have commonly had difficulties when wavefronts become multivalued (see Fig. 2). This occurs when rays cross or caustics appear, phenomena automatically handled when using a Lagrangian approach. The main issue is in avoiding the viscosity solution, which does not allow for multivalued wavefronts, of the Hamilton–Jacobi PDE. Much research in the past few years has been dedicated to solving this problem (see, e.g., [1, 5–7, 9, 14–17, 30, 36]), with mixed results. For a more detailed review of current methods, see [16]. We will also comment more on the approach taken in [16] since we have used the same basic setup as that study.

We present here a level set approach for following the representation of the wavefront in a higher dimensional reduced phase space. In fact, ray tracing in geometric optics is inherently higher dimensional. This phenomenon is not uncommon, for example, in radar scattering, where we are interested in not one wave scattering off an object but many waves, thus leading to higher dimensions. Standard phase space is the set of (x, p, t) , where $x \in \mathbf{R}^n$, $p \in \mathbf{R}^n$, $t \in \mathbf{R}^+$, which are linked to the problem through the Hamiltonian function




FIG. 2. Example of a multivalued and nonsmooth wavefront.

$H(x, p, t)$. Here, x represents the point in space, p the local phase direction, and t time. Note that in an isotropic medium, which we concentrate on in this paper, p is exactly the local ray direction. The wavefronts can be represented in phase space as the set of bicharacteristic strips Λ , a Lagrangian submanifold of codimension n (see, e.g., [3, 4, 13, 22]), with the projection of Λ to spatial space and time, $\mathbf{R}^2 \times \mathbf{R}^+$, giving back the wavefronts. Note that Λ is smooth even though its projection viewed at a particular time may be a multivalued solution (see Fig. 3). Hence, we endeavor to solve for Λ , through the Liouville equations (see, e.g., [16]) and thus take care of the difficulties of caustics and multivalued solutions




FIG. 3. The bicharacteristic strip in reduced phase space associated with the wavefront in Fig. 2. Note that it is smooth and vertical above the nonsmooth points of the wavefront.

by working in phase space. More precisely, we work with reduced Liouville equations that solve the equivalent submanifold in a reduced phase space composed of (x, Θ, t) , where $\Theta \in \mathbf{S}^{n-1}$ and $\Theta = \frac{p}{|p|}$. The information on the magnitude of p is not lost, appearing instead in the guise of the local wave velocity c under the relation $|p| = \frac{1}{c}$ in our final form for the reduced Liouville equations. Numerically, to obtain well-resolved wavefronts, we endeavor to use the level set Eulerian approach (see [28]) presented in [10] for dealing with objects of high codimension to represent the submanifold of interest. Under this representation, the reduced Liouville equations translate to a system of partial differential equations. Thus, altogether, we introduce here a Eulerian and PDE approach operating in reduced phase space for ray tracing and constructing wavefronts that handles multivalued solutions and spatial resolution of the wavefronts automatically.

2. LEVEL SET FORMULATION OF REDUCED LIOUVILLE EQUATIONS

For purposes of simplicity and exposition, we restrict our attention to $n = 2$ and comment on extensions to higher dimensions later on. Thus reduced phase space, for a fixed time, can be written as the set of (x, θ) , where $x \in \mathbf{R}^2$ and $\theta \in [-\pi, \pi]$, with θ being the angle of p in polar coordinates. So for a fixed time, the representation of the wavefront, called the bicharacteristic strip, is a smooth curve in reduced phase space (see Figs. 3 and 4). This fits nicely with [10], which considered motions of higher codimensional objects, mainly curves in \mathbf{R}^3 . In this formulation, the curve is represented by the intersection of two surfaces, which in turn are represented by the zero level sets of two real-valued functions in \mathbf{R}^3 , called the level set functions. Denoting the level set functions by ϕ and ψ , the curve is thus the set of points where $\phi = \psi = 0$ (see Fig. 5). Numerically, a uniform grid is placed in \mathbf{R}^3 and values of ϕ and ψ are given at the grid points. Periodic boundary conditions are also enforced to equate $\mathbf{R}^2 \times [-\pi, \pi]$ with \mathbf{R}^3 . Actually, we do not need a uniform grid in all of \mathbf{R}^3 since only the curve is of interest. We comment on this and its effect on efficiency later




FIG. 4. Curve in reduced phase space and its projection. Here θ is scaled to lie in $[-1, 1]$.




FIG. 5. The first picture shows the zero level set surfaces of two level set functions. The second picture shows the curve of intersection of those surfaces. The final picture shows the wavefront described by that curve. Here θ is scaled to lie in $[-1, 1]$.

on. This provides the desired Eulerian framework. In addition, even some methods using fixed Eulerian grids need interpolation to accurately represent the wavefront (see, e.g., [17, 30]); however, in our method, the representation and the PDE-based framework provides a self-interpolation property that automatically handles this. See [27] for more on level set methods.

By moving the bicharacteristic strip, we also move its projection, i.e., the wavefront. The velocity field for this motion, which does not depend on time, comes from the reduced Liouville equations of reduced phase space (see, e.g., [16]) and has the form

$$v(x, \theta) = \begin{pmatrix} c(x) \cos \theta \\ c(x) \sin \theta \\ c_{x_1}(x) \sin \theta - c_{x_2}(x) \cos \theta \end{pmatrix},$$

where c is the local wave velocity, determined from the index of refraction (see Fig. 6). This follows from the Hamiltonian system for ray tracing and agrees with the characteristics of the eikonal equation. Furthermore, it can be derived directly by a high-frequency approximation of the scaled Wigner distribution [44] (see [31]). Thus, adhering to the motion laws of [10],




FIG. 6. The wavefront, which is a circle, is shown on the left along with the directions of the velocity for index of refraction 1. The associated curve and directions in reduced phase space are shown on the right. As usual, θ is scaled to lie in $[-1, 1]$.

the system of evolution equations for the time-dependent level set functions under this velocity field is

$$\begin{aligned} \phi_t + v \cdot \nabla \phi &= 0, \\ \psi_t + v \cdot \nabla \psi &= 0, \end{aligned}$$

which form the level set formulation of the reduced Liouville equations. By solving this system of transport equations up to a certain time, we get the bicharacteristic strip at that time from the intersection of the zero level sets of ϕ and ψ . Hence the problem of constructing wavefronts becomes one of solving a system of partial differential equations in \mathbf{R}^3 . Note that if c is negative, the wavefront moves inward instead of outward. For more on the wave, eikonal, and Liouville equations, see the Appendix.

We note that the approach introduced in [16] also incorporates a Eulerian framework to deal with bicharacteristic strips, flowing under the reduced Liouville equations, in reduced phase space. For the two-dimensional problem, the method considers the projections, forming curves in \mathbf{R}^2 , of the desired curve in \mathbf{R}^3 from different directions. Then the segment projection method [37, 38] is used to represent the projected curves by viewing them as unions of graphs over different axes. These can then be evolved under Eulerian grids over the axes and recombined into the projected curves, which can in turn be recombined into the desired curve in \mathbf{R}^3 . Thus intricate data structures and techniques are needed to piece together the generally increasing number of different graphs to form a projected curve and different projected curves to form the desired curve, especially in the presence of numerical errors and in higher dimensions. However, we do note that aside from these considerations, the segment projection method is extremely efficient, as it works directly on points on the object of interest.

Note that another approach that has been implemented for handling motions of higher codimensional objects in a Eulerian framework is to simply use one level set function whose zero level set gives the curve in \mathbf{R}^3 (see, e.g., [2, 20, 26]). There is no need to worry about

thickening occurring during merging or breaking; this will not happen in a geometric optics problem, as the curve remains smooth for all time. However the zero level set is not stable under perturbations, resulting in a loss of accuracy in determining the location of the curve, which furthermore is usually modeled by a thin tube.

Finally, postdating our work, [19] also used the Liouville equation in a related problem. Theirs involved only one level set function, though it is different from the one level set function approach mentioned above. Thus the complexity of each update step is $\mathcal{O}(N)$ times ours for two-dimensional geometric optics with an $N \times N$ grid and $\mathcal{O}(N^2)$ times ours in three dimensions.

3. NUMERICALLY REALIZING THE FLOW

The evolution equations for the level set functions are decoupled transport equations and thus each one can be solved separately up to the desired time. For the spatial derivatives, we can use such high-order finite difference schemes for Hamilton–Jacobi equations as fifth-order WENO–Godunov (see [23]), made possible by the underlying uniform grid. For the time derivative, we can use, for example, third-order TVD–RK (see [33]) or the fourth-order SSP–RK of [34]. Iterating with these discretizations up to the desired time, regardless of whether the wavefronts are multivalued and without worrying about spatial resolution, gives the desired level set functions. The Courant–Friedrichs–Levy (CFL) condition on the allowable time step for stability in this case relates the time step to the spatial step of the grid in reduced phase space and the maximum velocity. For these evolution equations, we can take

$$\Delta t < C \frac{\Delta x}{\max \sqrt{|c|^2 + |\nabla c|^2}},$$

where Δt is the time step, Δx is the spatial step for a uniform grid, and C is a constant depending on the differencing in space and time used. Comparing this with the CFL condition for the eikonal equation solved using the same differencing techniques,

$$\Delta t < C \frac{\Delta x}{\max |c|},$$

shows that our approach requires a more restrictive time step. We comment more on this later, when dealing with discontinuous c , where the difference is more keenly felt. To plot the final curve and its projection from the final level set functions, a simplification of the algorithm in [25] involving interpolation can be used (see [10]). Finally, we on occasion scale θ to lie in $[-1, 1]$ via $\theta \mapsto \frac{\theta}{\pi}$ and use $c < 0$ for backward directions of motion in computations and diagrams.

In some cases, though, due to the appearance of what resemble kinks in the level set functions which may affect the PDE solvers and the plotting routine, reinitialization steps are used for regularization. These regularizations provide a nicer, more well-behaved form for the level set functions and theoretically do not move the location of the intersection of the zero level sets. The form we desire is signed distance functions for ϕ and ψ with the zero level sets orthogonal to each other at their points of intersection (see [10]). The signed distance representation for the level set functions spaces out their level set surfaces and the orthogonality requirement prevents the degeneration of the level set surfaces of ϕ being near

parallel to the level set surfaces of ψ . These considerations are important, for example, for accurately locating the position of the curve of interest for purposes of plotting. Producing signed distance functions is a well-studied topic and numerous algorithms can be used. One such method involves iterating

$$\begin{aligned}\tilde{\phi}_t + \text{sign}(\tilde{\phi}_0)(|\nabla\tilde{\phi}| - 1) &= 0, \\ \tilde{\psi}_t + \text{sign}(\tilde{\psi}_0)(|\nabla\tilde{\psi}| - 1) &= 0,\end{aligned}$$

to steady state at each time step, where $\tilde{\phi}_0$ and $\tilde{\psi}_0$ are ϕ and ψ , respectively, at that time step. The steady state solutions $\tilde{\phi}_\infty$ and $\tilde{\psi}_\infty$ comprise the new ϕ and ψ , respectively. This system of equations forces $|\nabla\phi| = 1$ and $|\nabla\psi| = 1$, the conditions for signed distance, without changing the zero level sets (see [35]). The advantage of this approach is the ability to generate very smooth, at least near the object of interest, and well-behaved approximations. Other methods that can be used for obtaining signed distance functions include fast marching and fast sweeping methods (see, e.g., [8, 12, 21, 32, 39, 40]).

Similarly, orthogonality can be achieved through fast marching and sweeping methods or by iterating

$$\begin{aligned}\tilde{\phi}_t + \text{sign}(\tilde{\psi}_0)\frac{\nabla\tilde{\psi}}{|\nabla\tilde{\psi}|} \cdot \nabla\tilde{\phi} &= 0, \\ \tilde{\psi}_t + \text{sign}(\tilde{\phi}_0)\frac{\nabla\tilde{\phi}}{|\nabla\tilde{\phi}|} \cdot \nabla\tilde{\psi} &= 0,\end{aligned}$$

to steady state at each time step, where $\tilde{\phi}_0$ and $\tilde{\psi}_0$ are ϕ and ψ , respectively, at that time step. The resulting steady state solutions $\tilde{\phi}_\infty$ and $\tilde{\psi}_\infty$ comprise the new ϕ and ψ , respectively. This system of equations forces the condition $\nabla\phi \cdot \nabla\psi = 0$ without changing the intersection of the zero level sets (see, e.g., [11]). Once again, the advantage of this approach is in smooth and well-behaved approximations.

Note that these reinitialization steps, especially orthogonalization, couple the derivation of ϕ with that of ψ . The order of enforcing these conditions is also important to obtain a better level set representation. Usually, we first obtain signed distance for ϕ , then orthogonality for ψ , then signed distance for ψ , then orthogonality for ϕ , with the process being repeated until steady state is reached. We however note that steady state does not actually need to be reached in practice for the reinitialization PDEs. They can simply be iterated a few times every few time steps to obtain a smoother representation. In our studies, we have found that two reinitialization steps every two time steps seems to work well. Finally, the reinitialization PDEs can be solved over the same uniform grid using high-order methods, such as third-order TVD–RK or fourth-order SSP–RK in time and fifth-order WENO–Godunov along with high-order central differencing in space with the CFL condition

$$\Delta t < C \Delta x,$$

where C is the same as previously used in the evolution equations when the same difference methods are applied, for both PDEs. We do observe in practice, however, that reinitialization moves the zero level sets that should remain fixed. This is due to the smoothing applied to the sign functions in the equations. Thus reinitialization usually reduces the order of accuracy of the whole algorithm to first order, even with all the high-order difference schemes used.

However, we do note that the reinitialization steps, though they slightly move the location of the object of interest, give high-order-accurate solutions of the reinitialization PDEs to this perturbed location. This especially means the result of reinitialization is a well-behaved set of level set functions representing an object that is perturbed a bit from its original location.

These reinitialization PDEs are needed, in fact, when speed and memory storage are a factor. As mentioned before, a static uniform grid is placed over reduced phase space and the reduced Liouville equations are solved over this grid. However, operating and storing values at all the grid points is unnecessary. Instead of considering all of reduced phase space, we can just consider a subdomain containing the curve, usually chosen to be the set of points of a certain approximate distance, proportional to Δx , away from the curve. Thus computations are made and values are stored only at grid points inside this subdomain. The only condition on the subdomain is the numerical condition that it contain enough grid points to accurately perform discretization of the various PDEs. This is the main idea behind the local level set method (see [29]). The reinitialization PDEs are important in this setting for creating the subdomain at each time step and to prevent the boundary of the subdomain from affecting the results. Using the level set functions after reinitialization, the subdomain, for example, can be chosen to be the set of points in reduced phase space where ϕ and ψ are both less than a constant times the grid step size. An alternative, which we actually use in an efficient-in-time local level set method for geometric optics, is to create the tube by using distance away from grid points neighboring the object of interest, which can be easily calculated. The tube can then be taken to be the set of points where this distance is less than a constant times the grid step size. This constant in both cases depends on the width of stencils needed in the high-order difference methods for the various PDEs. For example, the stencil for fifth-order WENO centered at a point requires three gridpoints on each side of the point. Thus, in order to use fifth-order WENO, a tube with a thickness radius of around $4\Delta x$ is necessary, as the object of interest may not lie on a grid point. We usually choose a tube of radius around $6\Delta x$ to ensure that further points around the object also will have adequate approximations. Of course the second approach does not require reinitialized level set functions but reinitialization is nonetheless important in another aspect. Essentially, the process passes information from the curve to the whole domain, thus effectively eliminating possible undesired effects the boundary of the tube may cause during the evolution of the level set functions. In fact, two reinitialization steps each time step in our local level set algorithm seems to work quite well. Combining all these elements, we obtain a local level set method that at worst performs $\mathcal{O}(N \log N)$ operations at each time step, if the grid size is $N \times N \times N$. Our efficient-in-time algorithm in fact may be better, using only $\mathcal{O}(N)$ operations, since it ignores efficiency in memory and thus avoids intricate data structures. Note that in [10], it was observed that a local algorithm may encounter difficulties during merging due to the numerous configurations the level set functions can take in the subdomain. However, mergings do not occur in this setting since the bicharacteristic strips remain smooth. Therefore, there is no hindrance here to a local and efficient level set method.

It is also not difficult to initialize the two level set functions to represent a given curve in \mathbf{R}^2 . For example, suppose the curve is smooth, strictly convex, and given by the parametrization $\tilde{x}(s)$. Then the level set functions can be chosen to be

$$\begin{aligned}\phi(x, \theta) &= x_1 - \tilde{x}_1(r(\theta)), \\ \psi(x, \theta) &= x_2 - \tilde{x}_2(r(\theta)),\end{aligned}$$

where $r(\theta)$ reparametrizes the curve \tilde{x} so that for any $\theta \in [-\pi, \pi]$, the normal of the curve at $s = r(\theta)$ has angle θ . An example is initialization of the ellipse given by the parametrization

$$\begin{aligned}x(s) &= b \cos s, \\y(s) &= \sin s,\end{aligned}$$

for $b > 0$ and $-\pi \leq s \leq \pi$. One direction for the normals of the ellipse thus takes the form $(\cos s, b \sin s)$ and so $r(\theta)$ satisfies

$$\tan s = b \tan r(\theta).$$

Therefore, ϕ and ψ can be initialized to be

$$\begin{aligned}\phi &= x - \frac{b^2}{\sqrt{b^2 + \tan^2 \theta}}, \\ \psi &= y - \frac{\tan \theta}{\sqrt{b^2 + \tan^2 \theta}}.\end{aligned}$$

Note that the forms of the above choice of level set functions only allow for strictly convex curves. To represent more general curves, we note that the two properties the bicharacteristic strips in reduced phase space need to satisfy are that the projection gives the desired curves and the θ value at each point gives the angle of the normal of the desired curves. Thus if the curve is given in two-dimensional space as the zero level set of a two-dimensional level set function $\tilde{\phi}(x)$, then notice that one direction for the normals of the zero level set takes the form $(\tilde{\phi}_{x_1}(x), \tilde{\phi}_{x_2}(x))$, the spatial gradient of $\tilde{\phi}(x)$. The level set functions can then be chosen to be

$$\begin{aligned}\phi(x, \theta) &= \tilde{\phi}(x), \\ \psi(x, \theta) &= \tilde{\phi}_{x_1}(x) \sin \theta - \tilde{\phi}_{x_2}(x) \cos \theta,\end{aligned}$$

with $\phi = 0$ satisfying the first requirement for the curve derived from $\phi = \psi = 0$ in reduced phase space, and $\psi = 0$ satisfying the second one. So $\phi = \psi = 0$ is the bicharacteristic strip corresponding to the given desired curve in \mathbf{R}^2 . An example is the initialization of the graph of a function given by $y = h(x)$. A level set function in \mathbf{R}^2 corresponding to this curve takes the form $\tilde{\phi} = y - h(x)$. Using this, we can initialize the level set functions ϕ and ψ as

$$\begin{aligned}\phi &= y - h(x), \\ \psi &= -\frac{1}{h'(x)} - \tan(\theta).\end{aligned}$$

This concept can be extended to cases in higher dimensions. For example, for $n = 3$, one level set function is used to satisfy the first requirement of projection and two more are used to satisfy the second requirement of angle of normals. Thus the level set functions can be easily initialized for a given wavefront in spatial space.

4. TESTS AND EXAMPLES

We would first like to remark that though we mostly use complicated high-order finite difference techniques involving fourth-order SSP–RK and fifth-order WENO for the PDEs involved here, their usage is completely due to our personal preferences and desire to incorporate the latest available PDE solvers. Other simpler schemes may be used in place of these. In fact, anyone with any kind of acceptable solver for Hamilton–Jacobi equations can effectively implement their version of our method.

We also note that the results in this section were mostly produced over $50 \times 50 \times 50$ grids for two-dimensional geometric optics since wavefronts appear to be well resolved by a grid of this size. Studies have been made, however, using grids of various other sizes and especially when the grid step size goes to zero.

We begin with examples under constant index of refraction. Exact solutions are abundant, so this case is only for testing purposes, where we can observe the algorithm’s handling of multivalued solutions as well as spatial resolution of the bicharacteristic strips and wavefronts. Note that the θ component of the velocity vector is zero. Also every point moves at the same speed. Thus, in our experiments, we observe the level set functions remain well-behaved during the flow and so reinitialization is not needed in the global level set approach, i.e., when working over the whole grid in reduced phase space. The model problem for constant index of refraction is a wavefront that is an initially shrinking ellipse. The wavefront becomes multivalued in different configurations at later times. Moreover, after a certain time, the initially shrinking wavefront starts to grow, which tests spatial resolution. Figure 7 shows the wavefronts constructed using our approach in the case of index of refraction 1. Figure 8 shows a selection of the same wavefronts at different times. Figure 9 shows the bicharacteristic strip in reduced phase space for one of the multivalued wavefronts. The vertical parts of the curve are exactly over the cusps of the wavefront. Note that the cusps of the wavefronts in all these cases have the requisite sharp profile. Figure 10 shows an




FIG. 7. Example of wavefronts plotted on the same graph at different times with the initial wavefront a horizontally elongated shrinking ellipse and an index of refraction 1.

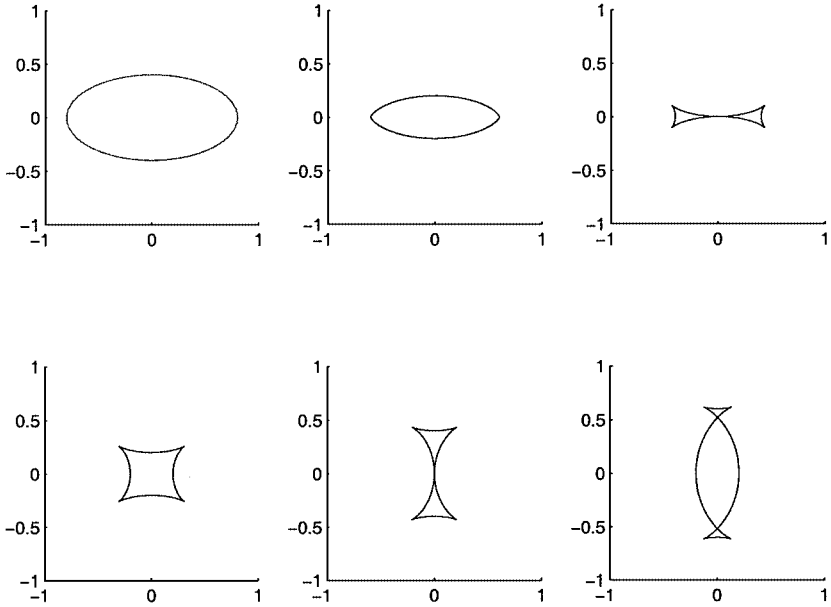


FIG. 8. A more detailed look at the evolution of the ellipse shown in Fig. 7, plotted at different times.

initially small ellipse expanding outward for an index of refraction 1. Note that the curve is well resolved at later time. Figure 11 reproduces the ray tracing solution shown in Figs. 1 and 2 in [5].

When the index of refraction is a general smooth function, reinitialization is needed to keep the level set functions well behaved during the flow. This is in part because the θ component of the velocity is no longer zero and points on the wavefronts in general do not move at the same speed. Model variable index of refraction cases include waveguide examples. Figures 12–14 are reproductions, using our algorithm, of examples run in [16]




FIG. 9. The figure on the left shows a wavefront with cusps from the evolution of an ellipse. The figure on the right shows the corresponding smooth bicharacteristic strip forming a curve in three-dimensional space. Here θ is scaled to lie in $[-1, 1]$ and c is taken to be negative.




FIG. 10. Initially small ellipse expanding and automatically being resolved.

and show the phenomenon with an initially straight wavefront at different times. Note that complicated multivalued solutions occur and remain well resolved. Also, Tables I and II show order-of-accuracy analyses for initially expanding ellipses in different mediums. No reinitialization was performed in these cases. The algorithm tested used SSP–RK of fourth order and WENO–Godunov of fifth order. Error was measured in the max norm at grid points neighboring the object of interest, where the numerically computed ϕ and ψ were compared to exact solutions. The fifth order results seen in the tables are commonly observed to occur, even in the presence of lower order time discretizations. Note that in terms of accuracy, the occurrence of multivalued wavefronts will not affect the accuracy of our algorithm since we deal with bicharacteristic strips which are always smooth.




FIG. 11. Reproduction under our method of the multivalued solution found in Figs. 1 and 2 of [5].




FIG. 12. An example of an initially straight wavefront moving to the right in a medium of variable index of refraction. In this case, the index of refraction is $\eta(x, y) = 1.5 - \frac{1}{\pi} \arctan(5(y - 1)^2 - 0.1(x - 0.5))$.

For our efficient-in-time local level set algorithm, we show order-of-accuracy and speed measurements in Tables III and IV. Table III shows results using SSP–RK of fourth order and WENO–Godunov of fifth order for both the evolution and reinitialization PDEs. Table IV shows results using SSP–RK of fourth order and WENO–Godunov of fifth order for the evolution equations and forward Euler with WENO–Godunov of fifth order for the reinitialization equations. Notice that reinitialization needs to be performed in local level set algorithms at each time step and hence only roughly first-order accuracy is exhibited. The error was calculated in the max norm by comparing points resulting from the second-order plotter to exact solutions, in this case a shrinking circle. Also notice that the speed




FIG. 13. Another example of an initially straight wavefront moving to the right in a medium of variable index of refraction. The index of refraction is given by $\eta(x, y) = 1 + e^{-y^2}$.




FIG. 14. Another example of an initially straight wavefront, lying on the y -axis, moving to the right in a medium of variable index of refraction. This is the same case as in Fig. 13, showing a wavefront at a larger travel time.

measurements show altogether roughly a $\mathcal{O}(N^2)$ computation time, or $\mathcal{O}(N)$ at each time step, since when the grid step size is halved, the computation time is around four times what it was.

The case of discontinuous index of refraction over the medium, for example piecewise constant local wave velocity, can also be handled. Numerically, this simply consists of approximating $c(x)$ by a smooth function and using the algorithm as before. However, one concern is in the restriction on the size of the time step that can be used due to stability requirements arising from the CFL number and condition. This restriction may seem severe because the time step depends inversely on the maximum of the velocity for the level set flow, which can be very large due to the terms containing derivatives of $c(x)$. In Fig. 15, we




FIG. 15. The figure on the left shows the results of a circular growing wavefront in a medium with discontinuous index of refraction using a standard Eulerian algorithm with spatial step $\Delta x = 0.01$. The figure on the right shows the results of our algorithm with $\Delta x = 0.04$ and obeying the CFL condition. In both cases, the wavefront travels from the left half plane with local wave velocity 2 to the right half plane with local wave velocity 1.

TABLE I

Order-of-Accuracy Analysis for the Global Algorithm without Reinitialization for Variable Local Wave Velocity $x^2 + y^2 + 1$ on an Initially Expanding Ellipse

| Grid size | Error | Order |
|-----------------------------|--------------------------|--------|
| $25 \times 25 \times 25$ | 0.0330223 | |
| $50 \times 50 \times 50$ | 0.00257825 | 3.6790 |
| $100 \times 100 \times 100$ | 0.000176317 | 3.8701 |
| $200 \times 200 \times 200$ | 5.82438×10^{-7} | 4.9199 |

Note. Error is calculated in the max norm.

TABLE II

Order-of-Accuracy Analysis for the Global Algorithm without Reinitialization for Variable Local Wave Velocity $\sin x \sin y + 1.5$ on an Initially Expanding Ellipse

| Grid size | Error | Order |
|-----------------------------|--------------------------|--------|
| $25 \times 25 \times 25$ | 0.00105497 | |
| $50 \times 50 \times 50$ | 6.27472×10^{-5} | 4.0715 |
| $100 \times 100 \times 100$ | 2.65032×10^{-6} | 4.5653 |
| $200 \times 200 \times 200$ | 8.46914×10^{-8} | 4.9678 |

Note. Error is calculated in the max norm.

TABLE III

Order-of-Accuracy and Computational-Time Analysis for the Local Algorithm for Variable Local Wave Velocity $0.5(\sin \sqrt{x^2 + y^2} + 1)$ on an Initially Shrinking Circle

| Grid size | Error | Order | Time |
|-----------------------------|-------------|--------|------|
| $50 \times 50 \times 50$ | 0.00107553 | | 120 |
| $100 \times 100 \times 100$ | 0.000543869 | 0.9837 | 466 |
| $200 \times 200 \times 200$ | 0.000198452 | 1.4545 | 2037 |

Note. SSP–RK of fourth order in time is used for the reinitialization PDEs. Error is calculated in the max norm and computational time is in seconds.

TABLE IV

Order-of-Accuracy and Computational-Time Analysis for the Local Algorithm for Variable Local Wave Velocity $0.5(\sin \sqrt{x^2 + y^2} + 1)$ on an Initially Shrinking Circle

| Grid size | Error | Order | Time |
|-----------------------------|-------------|--------|------|
| $50 \times 50 \times 50$ | 0.00105166 | | 41 |
| $100 \times 100 \times 100$ | 0.000416598 | 1.3359 | 159 |
| $200 \times 200 \times 200$ | 0.000203411 | 1.0343 | 690 |

Note. Forward Euler in time is used for the reinitialization PDEs. Error is calculated in the max norm and computational time is in seconds.




FIG. 16. The figure on the left shows the results of a circular growing wavefront in a medium with discontinuous index of refraction using a standard Eulerian algorithm with spatial step $\Delta x = 0.01$. The figure on the right shows the results of the CFL-violating version of our algorithm with $\Delta x = 0.02$. The prescribed local wave velocity is the same as in Fig. 15.

show results of refraction obeying the CFL condition. Note, however, that the velocity is only abnormally large at the discontinuities of $c(x)$, which generally form a set of measure zero in the medium. Thus, it is possible to simply consider violation of the CFL condition, i.e., using relatively large time steps, with hopes that the areas of discontinuity in $c(x)$ are too small to generate instabilities. For example, in the CFL condition, we can pretend that the maximum velocity is taken only at points away from the discontinuities. Thus the condition is valid everywhere except exactly at the discontinuities. Numerically, we note that we should not use a smooth approximation to $c(x)$ that enlarges the region of large velocity too much, as the CFL condition would be violated in too large a region, increasing the chances of instabilities forming. Also, the process of reinitialization itself is a regularizing effect. Thus without reinitialization, instabilities may develop, but with reinitialization, it seems we can violate the CFL condition without harm done to the final solution. In Fig. 16, we show results of refraction with large time steps that violate the CFL condition. Note that approximation of the discontinuity in the local wave velocity itself introduces error into the results. We have also run simulations checking our results with Snell's law in the case of straight line wavefronts hitting straight line interfaces that separate regions of different indices of refraction. We arrive at the correct angle of the wavefronts made at the interfaces in each case.

5. REFRACTION AND REFLECTION USING INTERFACE BOUNDARY CONDITIONS

Another approach to handling the transmission of rays in a medium of discontinuous index of refraction involves prescribing boundary conditions on the interface composed of the points lying on the discontinuity. Thus, instead of differencing across the interface, which can lead to large velocities and restrictive CFL numbers in the current scenario, it is possible to use the boundary conditions there instead. This approach can also be applied to reflection of rays off the boundary of an object. In the case of refraction, the interface boundary conditions can be derived from the behavior of rays passing through the

interface. This behavior is essentially defined by Snell's law. Thus, if α_I is the angle the incoming ray makes with the normal of the interface before transmission, α_T is the angle after transmission, c_I is the index of refraction at the interface before transmission, and c_T is the index after transmission, then these quantities satisfy

$$c_I \sin \alpha_T = c_T \sin \alpha_I.$$

This condition on rays can easily be translated to a condition on wavefronts and ultimately to conditions on the level set functions representing the wavefronts. The ghost fluid method [18] or other similar methods, such as the immersed interface method [24], can then be used to solve the level set system of PDEs for evolution with the above interface boundary conditions. This can remove the issue of restrictive CFL conditions that currently appears in our method.

Reflection can also be handled in a similar fashion. Unreflected and reflected rays at the interface satisfy the compatibility condition

$$\beta_I = \pi - \beta_R,$$

where β_I is the angle between the incoming ray and the normal at the interface before reflection and β_R is the angle after reflection. This translates into corresponding conditions for wavefronts at the interface. We introduce additional sets of level set functions to represent the wavefronts of reflected rays. The above interface boundary conditions couple these new level set functions to the original ones representing the wavefronts of unreflected rays. Then the ghost fluid method or other similar methods can be used to solve the resulting problem. Figure 17 shows results in the simple case of reflection off the walls of the grid boundary.




FIG. 17. Two examples of wavefronts reflecting off the walls of the grid boundaries.

In three dimensions, the interface boundary conditions take the form

$$C = -B + 2 \frac{(B \cdot A)A}{|A|^2},$$

where A is the normal vector to the interface at the point of reflection, B is the ray vector before reflection, and C is the ray vector after reflection. Similar conditions can also be derived in the three-dimensional case for refraction. A good implementation of these ideas is the subject of future work and would allow us to efficiently handle refraction and reflection of waves.

6. COMPUTING THE INTENSITY

As a postprocessing step, the intensity, or amplitude, of wavefronts can, in certain cases, be easily calculated from the results of the procedure. For example, for a curve parametrized by $(x(\theta, t), y(\theta, t), \theta)$ in reduced phase space, in a medium of constant index of refraction, the leading term A_0 of the intensity coming from the geometric optics approximation (see the Appendix) satisfies the equation

$$\frac{A_0(x, y, \theta, t)}{A_0(x, y, \theta, 0)} = \frac{\sqrt{x_\theta^2(\theta, 0) + y_\theta^2(\theta, 0)}}{\sqrt{x_\theta^2(\theta, t) + y_\theta^2(\theta, t)}}$$

(see, e.g., [16]). The quantity $x_\theta^2 + y_\theta^2$ can be expressed in our formulation as

$$x_\theta^2 + y_\theta^2 = \frac{J^2((\phi, \psi); (y, \theta)) + J^2((\phi, \psi); (x, \theta))}{J^2((\phi, \psi); (x, y))},$$

where J is the Jacobian, with

$$J^2((\phi, \psi); (y, \theta)) = \begin{vmatrix} \phi_y & \phi_\theta \\ \psi_y & \psi_\theta \end{vmatrix},$$

and similarly for $J((\phi, \psi); (x, \theta))$ and $J((\phi, \psi); (x, y))$. This is an easy and passive calculation to perform since A_0 just depends on the calculated values of ϕ and ψ . With ϕ and ψ , the location of the curve in reduced phase space can be found. For each point of the curve, the original point from whence it came can be determined since, in this case, it will be the point of the original curve with the same θ value. Thus, the intensity and Jacobian at the points of the original curve, as well as the Jacobian at the points of the current curve, can be computed using finite differences and interpolation. These inserted into the above equation then give the intensity on the desired curve at a given time.

In general, the intensity satisfies the equation

$$A_0^2 = \int w dp,$$

where $w = w(x, p, t)$ is the density (see, e.g., [16]). As usual, w can be rewritten in reduced form by replacing p , the ray vector, with $\frac{p}{|p|}$ and hence propagates according to the same

Liouville-type PDE that evolves the level set functions. Calculating the intensity is thus no longer a passive operation, requiring the density to be propagated along with curve location.

7. EXTENSIONS TO HIGHER DIMENSIONS AND MORE GENERAL HAMILTONIANS

Our approach can also be easily extended to higher dimensions, though computational speed and efficiency in memory storage start to become major issues. However, the physically relevant case $n = 3$ is still tractable. In this case, for fixed time, reduced phase space is five dimensional, $\mathbf{R}^3 \times \mathbf{S}^2$, and the bicharacteristic strips are codimension 3 manifolds. Even with these high dimensions and codimensions, the ideas and techniques in [10] still apply. We can use three level set functions ϕ, ψ , and η , mapping reduced phase space to \mathbf{R} , whose zero level sets intersect exactly at the bicharacteristic strip. Spherical coordinates can be used to represent points in \mathbf{S}^2 . The associated reduced Liouville equations, a system of three PDEs, can then be derived for these level set functions. They take the form

$$u_t + c(\cos \theta_1 \cos \theta_2 u_x + \cos \theta_1 \sin \theta_2 u_y + \sin \theta_1 u_z) + (c_x \sin \theta_1 \cos \theta_2 + c_y \sin \theta_1 \sin \theta_2 - c_z \cos \theta_1) u_{\theta_1} + \frac{(c_x \sin \theta_2 - c_y \cos \theta_2)}{\cos \theta_1} u_{\theta_2} = 0,$$

where $-\frac{\pi}{2} \leq \theta_1 \leq \frac{\pi}{2}$ and $-\pi \leq \theta_2 \leq \pi$ come from the spherical coordinates representation of p ,

$$p = |p|(\cos \theta_1 \cos \theta_2, \cos \theta_1 \sin \theta_2, \sin \theta_1).$$

Note that this is simply the system of transport evolution equations under the velocity

$$v(x, \theta_1, \theta_2) = \begin{pmatrix} c \cos \theta_1 \cos \theta_2 \\ c \cos \theta_1 \sin \theta_2 \\ c \sin \theta_1 \\ c_x \sin \theta_1 \cos \theta_2 + c_y \sin \theta_1 \sin \theta_2 - c_z \cos \theta_1 \\ \frac{(c_x \sin \theta_2 - c_y \cos \theta_2)}{\cos \theta_1} \end{pmatrix}.$$

As usual, we solve these equations over a static uniform grid in reduced phase space, thus simplifying high-order differencing methods and providing for automatic spatial resolution. Note that there may be some trouble at $\theta_1 = \pm \frac{\pi}{2}$, the north and south poles, since $\cos \theta_1 = 0$ there. This does not affect the constant index of refraction case since $|\nabla c| = 0$, canceling out the bad terms; however, for the more general case, the option of working in full phase space needs to be considered. Also, the corresponding CFL condition takes the form

$$\Delta t < C \frac{\Delta x}{\max \sqrt{|c|^2 + |\nabla c|^2}},$$

as in the two-dimensional case. The only real issue is efficiency, which can mostly be handled with a local level set method. The overhead introduced by the grid is not irretrievably high

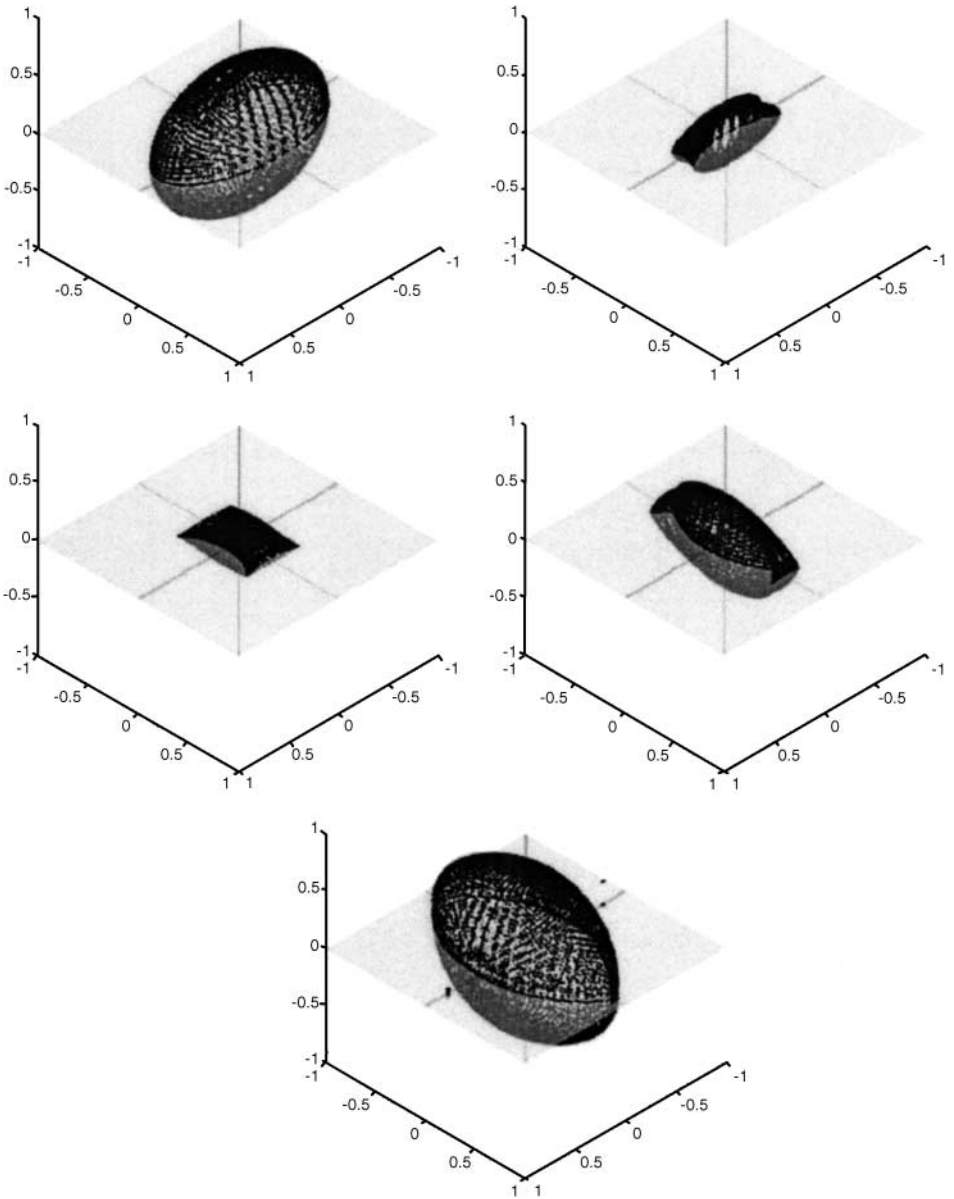


FIG. 18. Preliminary result of wavefronts in three dimensions over a coarse grid. This is an example, going from left to right, of an initially shrinking ellipse. Multivalued wavefronts seem to emerge. A better plotter is also needed for three-dimensional wavefronts.

for the case $n = 3$. Finally, in the reinitialization steps, creating signed distance functions carries through as before. Orthogonalization, however, would probably involve PDEs to enforce orthogonality of the level sets of ϕ to those of ψ and the level sets of η to both those of ϕ and ψ .

The full program for this case has not been completed; however, we have preliminary results. Figures 18 and 19 show an initially shrinking ellipse produced over a coarse $20 \times 20 \times 20 \times 20 \times 20$ grid. The corresponding initial level set functions can be taken




FIG. 19. Different views, from left to right, of the last three pictures in Fig. 18.

to be

$$\begin{aligned} \phi &= x - \frac{\alpha^2 \cos \theta_1 \cos \theta_2}{\sqrt{\gamma^2 \sin^2 \theta_1 + \cos^2 \theta_1 (\alpha^2 \cos^2 \theta_2 + \beta^2 \sin^2 \theta_2)}}, \\ \psi &= y - \frac{\beta^2 \cos \theta_1 \sin \theta_2}{\sqrt{\gamma^2 \sin^2 \theta_1 + \cos^2 \theta_1 (\alpha^2 \cos^2 \theta_2 + \beta^2 \sin^2 \theta_2)}}, \\ \eta &= z - \frac{\gamma^2 \sin \theta_1}{\sqrt{\gamma^2 \sin^2 \theta_1 + \cos^2 \theta_1 (\alpha^2 \cos^2 \theta_2 + \beta^2 \sin^2 \theta_2)}}, \end{aligned}$$

where $\alpha, \beta, \gamma > 0$ form the lengths of the principal axes of the ellipse. Note that the wavefronts become multivalued in the course of the flow. Also, the wavefronts at a certain time begin to expand. Both issues are properly handled using our approach. The plotter we use simply represents the wavefront as a cloud of points, as plotting routines that produce more aesthetically pleasing results, such as the one we used for two-dimensional geometric optics, are harder to implement in high dimensions. All this is due to the fact that the bicharacteristic

strips are only implicitly defined and so a postprocessing step, usually involving interpolation, is needed to determine their actual location. To plot the wavefront at a fixed time, we start with an initial set of points close to the bicharacteristic strip. This can simply be the grid points that are a certain distance, in some norm, away from the bicharacteristic strip. We then flow these points to the location of the bicharacteristic strip by minimizing $\phi^2 + \psi^2 + \eta^2$. Thus, in steady state, we get a cloud of points representing the bicharacteristic strip in reduced phase space. The points are then projected to spatial space to form a cloud of points representing the wavefront. Note, it may be possible to run a triangulated surface through the cloud of points when in reduced phase space using algorithms based on Delauney triangulation and such. Then the projection is also an interpolated surface instead of a set of discrete points. It may also be possible to use advanced algorithms for implicit surfaces such as in [45] to create a better plotter.

For even higher dimensions n , the algorithm follows the same guidelines and procedures. For a fixed time, reduced phase space takes the form $\mathbf{R}^n \times \mathbf{S}^{n-1}$, which has dimension $2n - 1$, with bicharacteristic strips of dimension $n - 1$. Thus the codimension of the bicharacteristic strips is n . This can be handled by introducing the vector-valued level set function ϕ that depends on time and maps the space \mathbf{R}^{2n-1} to \mathbf{R}^n . So ϕ takes the form

$$\phi = (\phi_1, \phi_2, \dots, \phi_n),$$

where ϕ_i is a scalar level set function for all i . The zeros of ϕ for a fixed time can then represent the bicharacteristic strips at that time and PDEs for the evolution of ϕ can be derived from equations governing the position and ray vectors of the wavefronts. Numerically, uniform grids can be placed over \mathbf{R}^{2n-1} and the evolution PDEs can be discretized using Hamilton–Jacobi solvers that work in high dimensions, such as ENO (essentially nonoscillatory) [33] or WENO, along with TVD–RK or SSP–RK schemes. Equations for reinitialization can also be produced and used in these higher dimensional spaces. The main problem with this approach in higher dimensions is that the efficiency of the algorithm is inversely related to the dimension of reduced phase space due to the grid-based numerics and so calculations for large n may be impractical.

Geometric optics can also be thought of as the model problem for first-order Hamilton–Jacobi equations with convex Hamiltonians. We can, in fact, extend our method to handle more general Hamilton–Jacobi equations,

$$\tilde{\phi}_t + H(x, \nabla \tilde{\phi}) = 0,$$

and for $\nabla \tilde{\phi} = p$,

$$\begin{aligned} \frac{dx}{dt} &= H_p, \\ \frac{dp}{dt} &= -H_x. \end{aligned}$$

The corresponding Liouville equation takes the form

$$u_t + H_p \cdot \nabla_x u - H_x \cdot \nabla_p u = 0,$$




FIG. 20. Preliminary results for the evolution of an open curve.

as usual. Using this in our level set representation, it is possible to simulate ray tracing for more general Hamilton–Jacobi equations.

Finally, problems dealing with open curves can now be attempted using our approach. Considering the projection of a curve in $\mathbf{R}^2 \times [-\pi, \pi]$ down to \mathbf{R}^2 that is not 2π periodic in the θ direction provides a method for smoothly representing an open and possibly multi-valued curve in two dimensions. Further representation of the curve in $\mathbf{R}^2 \times [-\pi, \pi]$ using our two level set function approach gives a Eulerian framework and the ability to handle geometric quantities with ease. In Fig. 20, we show preliminary results for the evolution of an open curve that displays multivalued solutions.

8. CONCLUSION

Thus we have introduced a PDE-based algorithm for ray tracing in geometric optics that not only handles multivalued solutions but takes care of spatial resolution without need of interpolation. The basic framework may be as in [16] but the approach to obtaining a Eulerian method is different, borrowing from work on geometric motions of higher codimensional objects [10]. This approach is simple and easily implemented in two dimensions, with all the equations and ideas also extendable to three dimensions. Furthermore, we have been able to create efficient-in-time algorithms that increase the usability of our method and its application to realistic cases in the area.

APPENDIX: THE WAVE, EIKONAL, AND LIOUVILLE EQUATIONS

The basic wave equation takes the form

$$\frac{\partial^2 u}{\partial t^2}(x, t) = c(x)^2 \Delta u(x, t),$$

with $x \in \Omega \subset \mathbf{R}^n$, $t > 0$, and initial conditions

$$\begin{aligned} u(x, 0) &= u_0(x), \\ \frac{\partial u}{\partial t}(x, 0) &= u_1(x), \end{aligned}$$

along with certain boundary conditions on $\partial\Omega$. For high-frequency wave propagation, inserting the ansatz

$$u \approx e^{i\omega\tilde{\Phi}(x,t)} \sum_{j=0}^{\infty} A_j(x,t)(i\omega)^{-j}$$

into the wave equation leads to

$$\omega^2 \left(\left(\frac{\partial \tilde{\Phi}}{\partial t} \right)^2 - c^2 |\nabla \tilde{\Phi}|^2 \right) = 0$$

as the leading coefficient term. Thus

$$\frac{\partial \tilde{\Phi}}{\partial t} \pm c |\nabla \tilde{\Phi}| = 0,$$

with a choice of plus or minus. This is the eikonal equation. We actually want the multivalued solutions, as seen in Figs. 2 and 8, to these equations. Ray tracing essentially involves the method of characteristics on the eikonal equations, though it can be derived independently. In fact, the characteristics for

$$\tilde{\Phi}_t + c |\nabla \tilde{\Phi}| = 0,$$

where $\nabla \tilde{\Phi} = (p_1, p_2, \dots, p_n) = p$, are given by the system of ordinary differential equations

$$\begin{aligned} \frac{dx}{dt} &= c \frac{p}{|p|}, \\ \frac{dp}{dt} &= -|p| \nabla c. \end{aligned}$$

These equations give the velocity field under which the wavefronts move. It is also possible to derive from this a PDE that transports values of a function along these characteristic directions. Let $w(x, p, t)$ be a function, usually denoting the density, where $x = (x_1, x_2, \dots, x_n)$, $p = (p_1, p_2, \dots, p_n)$, and $t > 0$. Thus w is a function of $2n + 1$ independent variables. The PDE for w that transports its values along the characteristic directions then takes the form

$$w_t + c \frac{p}{|p|} \cdot \nabla_x w - |p| \nabla_x c \cdot \nabla_p w = 0.$$

This is the Liouville equation. As should be, the linear equation and characteristics are

$$\begin{aligned} \frac{dx}{dt} &= c \frac{p}{|p|}, \\ \frac{dp}{dt} &= -|p| \nabla c. \end{aligned}$$

along which w is constant.

ACKNOWLEDGMENTS

We thank the Institute for Pure and Applied Mathematics (IPAM) for its Geometrically Based Motions program held in the spring of 2001, which introduced us to and motivated us in this subject matter.

REFERENCES

1. R. Abgrall and J.-D. Benamou, Big ray tracing and Eikonal solver on unstructured grids: Application to the computation of a multivalued travel-time field in the Marmousi model, *Geophysics* **64**, 230 (1999).
2. L. Ambrosio and H. M. Soner, Level set approach to mean curvature flow in arbitrary codimension, *J. Differential Geometry* **43**, 693 (1996).
3. V. I. Arnol'd, *Catastrophe Theory* (Springer-Verlag, Berlin/New York, 1992).
4. V. I. Arnol'd, S. M. Gusein-Zade, and A. N. Varchenko, *Singularities of Differential Maps* (Birkhauser, Basel, 1986).
5. J.-B. Benamou, An Eulerian numerical method for geometric optics, *COCV Proc.*, in press.
6. J.-D. Benamou, Big ray tracing: Multi-valued travel time field computation using viscosity solutions of the Eikonal equation, *J. Comput. Phys.* **128**, 463 (1996).
7. J.-D. Benamou, Direct solution of multi-valued phase-space solutions for Hamilton–Jacobi equations, *Commun. Pure Appl. Math.* **52** (1999).
8. M. Boué and P. Dupuis, Markov chain approximations for deterministic control problems with Affine dynamics and quadratic cost in the control, *SIAM J. Numer. Anal.* **36**(3), 667 (1999).
9. Y. Brenier and L. Corrias, A kinetic formulation for multi-branch entropy solutions of scalar conservation laws, in *Ann. IHP Analyse Non-Lineaire*, Vol. 15, No. 2, pp. 169–190, 1998.
10. P. Burchard, L.-T. Cheng, B. Merriman, and S. Osher, Motion of curves in three spatial dimensions using a level set approach, *J. Comput. Phys.* **170**, 720 (2001).
11. S. Chen, B. Merriman, S. Osher, and P. Smereka, A simple level set method for solving Stefan problems, *J. Comput. Phys.* **135**, 8 (1997).
12. P. Danielsson, Euclidean distance mapping, *Comput. Graphics Image Process.* **14**, 227 (1980).
13. J. J. Duistermaat, Oscillatory integrals, Lagrange immersions and unfolding of singularities, *Commun. Pure Appl. Math.* **27**, 207 (1974).
14. B. Engquist, E. Fatemi, and S. Osher, Numerical resolution of the high frequency asymptotic expansion of the scalar wave equation, *J. Comput. Phys.* **120**, 145 (1995).
15. B. Engquist and O. Runborg, *Multi-Phase Computation in Geometrical Optics*, Tech Report (Nada KTH, 1995).
16. B. Engquist, O. Runborg, and A.-K. Tornberg, *High Frequency Wave Propagation by the Segment Projection Method*, CAM Report 01-13 (UCLA, 2001).
17. J. Steinhoff, M. Fan, and L. Wang, A new Eulerian method for the computation of propagating short acoustic and electromagnetic pulses, *J. Comput. Phys.* **157**, 683 (2000).
18. R. Fedkiw, T. Aslam, B. Merriman, and S. Osher, A non-oscillatory Eulerian approach to interfaces in multimaterial flows (the ghost fluid method), *J. Comput. Phys.* **152**, 457 (1999).
19. S. Fomel and J. A. Sethian, *Fast Phase Space Computation of Multiple Arrivals*, Technical Report LBNL-48976 (LBL, 2001).
20. E. De Giorgi, *Barriers, Boundaries, Motion of Manifolds* (Conferenza al Dipartimento di Matematica dell'Università di Pavia, 1994).
21. J. Helmsen, E. Puckett, P. Colella, and M. Dorr, Two new methods for simulating photolithography development in 3D, *Proc. SPIE* **2726**, 253 (1996).
22. S. Izumiya, The theory of Legendrian unfoldings and first order differential equations, *Proc. Soc. Edinburgh* **123A**, 517 (1993).
23. G. S. Jiang and D. Peng, Weighted ENO schemes for Hamilton Jacobi equations, *SIAM J. Sci. Comput.* **21**(6), 2126 (2000).

24. Z. Li, *The Immersed Interface Method—A Numerical Approach to Partial Differential Equations with Interfaces*, Ph.D. thesis (University of Washington, 1994).
25. W. E. Lorensen and H. E. Cline, Marching cubes: A high resolution 3D surface reconstruction algorithm, *Comput. Graphics* **21**(4), 163 (1987).
26. L. M. Lorigo, O. Faugeras, W. E. L. Grimson, R. Keriven, R. Kikinis, and C. F. Westin, Co-dimension 2 geodesic active contours for MRA segmentation, in *International Conference on Information Processing in Medical Imaging*, June/July 1999 (Springer-Verlag, Berlin/New York, 1999) pp. 126–139.
27. S. Osher and R. Fedkiw, Level set methods: An overview and some recent results, *J. Comput. Phys.* **169**, 463 (2001).
28. S. Osher and J. A. Sethian, Fronts propagating with curvature dependent speed: Algorithms based on Hamilton–Jacobi formulations, *J. Comput. Phys.* **169**, 12 (1988).
29. D. Peng, B. Merriman, S. Osher, H. Zhao, and M. Kang, A PDE-based fast local level set method, *J. Comput. Phys.* **155**, 410 (1999).
30. S. J. Ruuth, B. Merriman, and S. Osher, Fixed grid method for capturing the motion of self-intersecting interfaces and related PDEs, *J. Comput. Phys.* **163**, 1 (2000).
31. L. Ryzhik, G. Papanicolaou, and J. Keller, Transport equations for elastic and other waves in random media, *Wave Motion* **24**(4), 327 (1996).
32. J. A. Sethian, Fast marching level set methods for three dimensional photolithography development, *Proc. SPIE* **2726**, 261 (1996).
33. C. W. Shu and S. Osher, Efficient implementation of essentially nonoscillatory shock-capturing schemes, *J. Comput. Phys.* **77**, 439 (1988).
34. R. J. Spiteri and S. J. Ruuth, A new class of optimal high-order strong-stability-preserving time discretization methods, *SINUM*, in press.
35. M. Sussman, P. Smereka, and S. Osher, A level set method for computing solutions to incompressible two-phase flow, *J. Comput. Phys.* **114**, 146 (1994).
36. W. Symes, *A Slowness Matching Finite Difference Method for Travel times beyond Transmission Caustics*, The Rice Inversion Project Annual Report (Rice University, 1996).
37. A.-K. Tornberg, *Interface Tracking Methods with Applications to Multiphase Flow*, Ph.D. thesis. (Department of Numerical Analysis and Computer Science, Royal Institute of Technology (KTH), Stockholm, 2000).
38. A.-K. Tornberg and B. Engquist, Interface tracking in two-phase flows, in *Multifield Problems in Solid and Fluid Mechanics* (Springer-Verlag, Berlin/New York, 2000).
39. Y.-H. Tsai, L.-T. Cheng, S. Osher, and H. K. Zhao, Fast sweeping algorithms for a class of Hamilton–Jacobi equations, UCLA CAM Report 01–27, 2001, submitted to SINUM.
40. J. N. Tsitsiklis, Efficient algorithms for globally optimal trajectories, *IEEE Trans. Autom. Control* **50**, 1528 (1995).
41. V. Vinje, E. Iversen, K. Astebol, and H. Gjøystdal, Estimation of multivalued arrivals in 3D models using wavefront construction—Part I, *Geophys. Prospect.* **44**, 819 (1996).
42. V. Vinje, E. Iversen, K. Astebol, and H. Gjøystdal, Part II: Tracing and interpolation, *Geophys. Prospect.* **44**, 843 (1996).
43. V. Vinje, E. Iversen, and H. Gjøystdal, Traveltime and amplitude estimation using wavefront construction, *Geophysics* **58**(8), 1157 (1993).
44. E. Wigner, On the quantum correction for thermodynamic equilibrium, *Phys. Rev.* **40**, 749 (1932).
45. H. K. Zhao, S. Osher, B. Merriman, and M. Kang, Implicit and nonparametric shape reconstruction from unorganized points using variational level set method, *Comput. Vision Image Understand.* **80**, 295 (2000).