

# JET MARCHING METHODS FOR SOLVING THE EIKONAL EQUATION

SAMUEL F. POTTER AND MARIA K. CAMERON

**Abstract.** We develop a family of compact high-order semi-Lagrangian label-setting methods for solving the eikonal equation. These solvers march the total 1-jet of the eikonal, and use Hermite interpolation to approximate the eikonal and parametrize characteristics locally for each semi-Lagrangian update. We describe solvers on unstructured meshes in any dimension, and conduct numerical experiments on regular grids in two dimensions. Our results show that these solvers yield at least second-order convergence, and, in special cases such as a linear speed of sound, third-order of convergence for both the eikonal and its gradient. We additionally show how to march the second partials of the eikonal using cell-based interpolants. Second derivative information computed this way is frequently second-order accurate, suitable for locally solving the transport equation. This provides a means of marching the prefactor coming from the WKB approximation of the Helmholtz equation. These solvers are designed specifically for computing a high-frequency approximation of the Helmholtz equation in a complicated environment with a slowly varying speed of sound, and, to the best of our knowledge, are the first solvers with these properties. We provide a link to a package online providing the solvers, and from which the results of this paper can be reproduced easily.

**Key words.** eikonal equation, high-order solver, semi-Lagrangian solver, Hermite interpolation, direct solver, marching, Helmholtz equation, geometric spreading

**AMS subject classifications.** 65N99, 65Y20, 49M99

**1. Introduction.** Our goal is to develop a family of high-order semi-Lagrangian eikonal solvers which use compact stencils. This is motivated by problems in high-frequency room acoustics, although the eikonal equation arises in a tremendous variety of modeling problems [40].

In multimedia, virtual reality, and video games, precomputing room impulse responses (RIRs) or transfer functions (RTFs) enables convincing spatialized audio, in combination with binaural or surround sound formats. Such an approach, usually referred to as *numerical acoustics*, involves computing pairs of RIRs by placing probes at different locations in a voxelized domain, numerically solving the acoustic wave equation, and capturing salient perceptual parameters throughout the domain using a streaming encoder [33, 34]. These parameters are later decoded using signal processing techniques in real time as the listener moves throughout the virtual environment. Assuming that the encoded parameters can comfortably fit into memory, a drawback of this approach is that the complexity of the simulation depends intrinsically on the highest frequency simulated. In practice, simulations top out at around 1 kHz. The hearing range of humans is roughly 20 Hz to 20 kHz, which requires these methods to either implicitly or explicitly extrapolate the bandlimited transfer functions to the full audible spectrum.

An established alternative to this approach is *geometric acoustics*, where methods based on raytracing are used [35]. Contrary to methods familiar from computer graphics, the focus of geometric acoustics is different. Acoustic waves are mechanical and have macroscopic wavelengths. This means that subsurface scattering, typically modeled using BRDFs in raytracing for computer graphics [27], is less relevant, and is limited to modeling macroscopic scattering from small geometric features, since reflections from flat surfaces are specular in nature. What’s more, accurately modeling diffraction effects is crucial [36]: e.g., we can hear a sound source occluded by an obstacle, but we can’t see it. A variety of other geometric-acoustic methods exist beyond raytracing. Examples include the image source method [2] and frustum tracing [13].

Geometric acoustics and optics both assume a solution to the wave equation based on an asymptotic high-frequency (WKB) approximation to the Helmholtz equation [29]. In this approximation, the eikonal plays the role of a spatially varying phase function, whose level sets describe propagating wavefronts. The prefactor of this approximation describes the amplitude of these wavefronts. The WKB approximation assumes a ray of “infinite frequency”, suitable for optics, since the effects of diffraction are limited. A variety of mechanisms for augmenting this approximation with frequency-dependent diffraction effects have been proposed, the most successful of which is Keller’s *geometric theory of diffraction* [21] (including the later *uniform theory of diffraction* [23]).

The complete geometric acoustic field of multiply reflected and diffracted rays can be parametrized by repeatedly solving the eikonal equation, using boundary conditions derived from the WKB approximation to patch together successive fields. A related approach is Benamou’s *big raytracing* (BRT) [5, 6]. This approach requires one to be able to accurately solve the transport equation describing the amplitude, e.g. using paraxial raytracing [29]. In order to do this, the first and second order partial derivatives of the eikonal must be computed. High-order accurate iterative schemes for solving the eikonal equation exist [50, 48, 24], but their performance deteriorates in the presence of complicated obstacles. Direct solvers for the eikonal equation allow one to locally parametrize the characteristics (rays) of the eikonal equation, which puts one in a position to simultaneously march the amplitude. This enables work-efficient algorithms, critical if a large number of eikonal problems must be solved.

Benamou’s line of research related to BRT seems to have stalled due to difficulties faced with caustics [7]. This is reasonable considering that the intended use was seismic modeling, where the eikonal equation is used to model first arrival times of  $P$ -waves. In this case, the speed of sound is extremely complicated, resulting in a large number of caustics [47]. On the other hand, in room acoustics, the speed of sound varies slowly. The main challenge is geometric: the domain is potentially filled with obstacles. This provides another motivation for compact stencils: such stencils can be adapted for use with unstructured meshes, and the sort of complicated boundary conditions that arise when using finite differences are avoided entirely. In this work, our goal is to develop the underlying approach to obtaining a compact higher-order semi-Lagrangian eikonal solver. In future work, this will be applied in the unstructured setting.

The solvers developed in this work are high-order, have optimally local/compact stencils, and are label-setting methods (much like Sethian’s *fast marching method* [39] or Tsitsiklis’s semi-Lagrangian algorithm for solving the eikonal equation [46]). Additionally, being semi-Lagrangian, they locally parametrize characteristics (acoustic rays), making them suitable for use with paraxial raytracing [29], the method of choice for locally computing the amplitude. To the best of our knowledge, these are the first eikonal solvers with this collection of properties.

We refer to our solvers as *jet marching methods* to reflect the fact that the key idea is marching the *jet* of the eikonal (the eikonal and its partial derivatives up to a particular order [43]) in a principled fashion. Sethian and Vladimirov developed a fast marching method that additionally marched the gradient of the eikonal in a short note, but did not prove convergence results or provide detailed numerical experiments [41]. Benamou and collaborators built on these ideas by exploiting information about the eikonal equation to obtain a compact upwind second-order finite difference scheme for solving the eikonal equation [8]. Related methods exist in the level set method community and are referred to as *gradient-augmented level set methods* or *jet*

schemes [25, 37].

In the rest of this work we lay out these methods, providing detailed numerical experiments. Our presentation is for unstructured grids in  $n$ -dimensions, while our numerical experiments were carried out in 2D. We plan to extend these solvers to structured and unstructured meshes in 3D and will report on these later in the context of room acoustics applications.

**1.1. Problem setup.** Let  $\Omega \subseteq \mathbb{R}^n$  be a domain, let  $\partial\Omega$  be its boundary, and let  $\Gamma \subseteq \Omega$ . The *eikonal equation* is a nonlinear first-order hyperbolic partial differential equation given by:

$$(1.1) \quad \begin{aligned} \|\nabla\tau(\mathbf{x})\| &= s(\mathbf{x}), & \mathbf{x} \in \Omega, \\ \tau(\mathbf{x}) &= g(\mathbf{x}), & \mathbf{x} \in \Gamma. \end{aligned}$$

Here,  $\tau : \Omega \rightarrow \mathbb{R}$  is the *eikonal*, a spatial phase function that encodes the first arrival time of a wavefront propagating with pointwise *slowness* specified by  $s : \Omega \rightarrow (0, \infty)$ , which can be thought of as an index of refraction. The function  $g : \Gamma \rightarrow \mathbb{R}$  specifies the boundary conditions, and is subject to certain compatibility conditions [9].

One way of arriving at the eikonal equation is by approximating the solution  $u$  of the *Helmholtz equation*:

$$(1.2) \quad \left(\Delta + \omega^2 s(\mathbf{x})^2\right)u(\mathbf{x}) = 0,$$

with the WKB ansatz:

$$(1.3) \quad u(\mathbf{x}) \sim \alpha(\mathbf{x})e^{i\omega\tau(\mathbf{x})},$$

where  $\omega$  is the frequency [29]. As  $\omega \rightarrow \infty$ , this asymptotic approximation is  $O(\omega^{-1})$  accurate. This is referred to as the *geometric optics* approximation [7]. The level sets of  $\tau$  denote the arrival times of bundles of rays, and the amplitude  $\alpha$ , which satisfies the transport equation:

$$(1.4) \quad \alpha(\mathbf{x})\Delta\tau(\mathbf{x}) + 2\nabla\tau(\mathbf{x})^\top \nabla\alpha(\mathbf{x}) = 0,$$

describes the attenuation of the amplitude of the wavefront due to the propagation and geometric spreading of rays. The characteristics (*rays*) of the eikonal equation satisfy the raytracing ODEs.

The solution of the eikonal equation is given by Fermat's principle:

$$(1.5) \quad \tau(\mathbf{x}) = \min_{\substack{\mathbf{y} \in \Gamma \\ \psi: [0,1] \rightarrow \Omega \\ \psi(0)=\mathbf{y}, \psi(1)=\mathbf{x}}} \left\{ \tau(\mathbf{y}) + \int_0^1 s(\psi(\sigma)) \|\psi'(\sigma)\| d\sigma \right\}.$$

Observe that this equation is recursive, suggesting a connection with dynamic programming and Bellman's principle of optimality. Indeed, the path  $\psi$  is a ray whose Lagrangian and Hamiltonian are:

$$(1.6) \quad \mathcal{H}(\mathbf{x}, \nabla\tau(\mathbf{x})) = \frac{\|\nabla\tau(\mathbf{x})\|^2 - s(\mathbf{x})^2}{2} = 0, \quad \mathcal{L}(\mathbf{x}, \dot{\mathbf{x}}) = s(\mathbf{x})\|\dot{\mathbf{x}}\|.$$

This provides the connection between the Eulerian perspective given by the eikonal equation, Fermat's principle, and the Lagrangian view provided by raytracing.

**2. Related work.** The quintessential numerical method for solving the eikonal equation is the *fast marching method* [38]. We discretize  $\Omega$  into a grid of nodes  $\Omega_h$ , where  $h > 0$  is the characteristic length scale of elements in  $\Omega_h$ . Let  $T : \Omega_h \rightarrow \mathbb{R}$  be the numerical eikonal. To compute  $T$ , equation (1.1) is discretized using first-order finite differences and the order in which individual values of  $T$  are relaxed is determined using a variation of Dijkstra’s algorithm for solving the single source shortest paths problem [39, 40]. If  $N = |\Omega_h|$ , then the fast marching method solves (1.1) in  $O(N \log N)$  with  $O(h \log \frac{1}{h})$  worst-case accuracy [51]. The logarithmic factor only appears when rarefaction fans are present: e.g., point source boundary data, or if the wavefront diffracts around a singular corner or edge. In these cases, full  $O(h)$  accuracy can be recovered by proper initialization near rarefaction fans, or by employing a variety of factoring schemes [17, 24, 31].

It is also possible to solve the eikonal equation using semi-Lagrangian numerical methods, in which the ansatz (1.5) is discretized and applied locally [46, 30]. For instance, at a point  $\hat{x} \in \Omega_h$ , we consider a neighborhood of points  $\text{nb}(\hat{x}) \subseteq \Omega_h$ , assume that  $\tau$  is fixed over the “surface” of this neighborhood, and approximate (1.5). As an example, if  $\text{nb}(\hat{x})$  consists of its  $2n$  nearest neighbors, if we linearly interpolate  $\tau$  over the facets of  $\text{conv}(\text{nb}(\hat{x}))$ , and discretize the integral in (1.5) using a right-hand rule, the resulting solver is equivalent to the fast marching method [42].

The Eulerian approach has generally been favored when developing higher-order solvers for the eikonal equation [50]. The eikonal equation is discretized using higher-order finite difference schemes and solved in the same manner as the fast marching method or using a variety of appropriate iterative schemes. Unfortunately, these approaches presuppose a regular grid and require wide stencils.

Our goal is to develop solvers for the eikonal equation that are *high order*, are *optimally local* (only use information from the nodes in  $\text{nb}(\hat{x})$  to update  $\hat{x}$ ), and are flexible enough to work on *unstructured meshes*. Using a semi-Lagrangian approach based on a high-order discretization of (1.5) allows us to do this.

This work was inspired by several lines of research. First, are *gradient-augmented level set methods* (or *jet schemes*) [25, 37]. Although developed for solving time-dependent advection problems, trying to map ideas from the *time-dependent to static* setting is natural, and presented an intriguing challenge. Second, the idea of using a semi-Lagrangian solver to construct a finite element solution to the eikonal equation incrementally was informative [9]; while the authors only constructed a first-order finite element approximation, attempting to push past this formulation to obtain a higher-order solver is a natural extension. Third, we were motivated by Chopp’s idea of building up piecewise bicubic interpolants locally while marching the eikonal [14]; indeed, Chopp’s work is mentioned in the original work on jet schemes in a similar capacity.

**3. The jet marching method.** Label-setting algorithms [12], such as the fast marching method, compute an approximation to  $\tau$  by marching a numerical approximation  $T : \Omega_h \rightarrow \mathbb{R}^n$  throughout the domain. The boundary data  $g$  is not always specified at the nodes of  $\Omega_h$ . Let  $\Gamma_h$  be a discrete approximation of  $\Gamma$ . Once  $T$  is computed at  $\Gamma_h \subseteq \Omega_h$  with sufficiently high accuracy, the solver begins to operate. To drive the solver, a set of states `{far, trial, valid}` is used for bookkeeping. We initially set:

$$(3.1) \quad \text{state}(\mathbf{x}) = \begin{cases} \text{trial}, & \text{if } \mathbf{x} \in \Gamma_h, \\ \text{far}, & \text{otherwise.} \end{cases}$$

The **trial** nodes are typically sorted by their  $T$  value into an array-based binary heap implementing a priority queue, although alternatives have been explored [18]. At each step of the iteration, the node  $\mathbf{x}$  with the minimum  $T$  value is removed from the heap,  $\text{state}(\mathbf{x})$  is set to **valid**, the **far** nodes in  $\text{nb}(\mathbf{x})$  have their state set to **trial**, and each **trial** node in  $\text{nb}(\mathbf{x})$  is subsequently updated. We have additionally provided a video online which shows the algorithm running [11].

From this, we can see that the value  $T(\mathbf{x})$  depends on the values of  $T$  at the nodes of a directed graph leading from  $\mathbf{x}$  back to  $\Gamma_h$ , noting that  $T(\mathbf{x})$  can—and in general does—depend on multiple nodes in  $\text{nb}(\mathbf{x})$ . This means that the error in  $T$  accumulates as the solution propagates downwind from  $\Gamma_h$ . We generally assume that the depth of the directed graph of updates connecting each  $\mathbf{x} \in \Omega_h$  to  $\Gamma_h$  is  $O(h^{-1})$ . The error due to each update comes from two sources: the running error accumulated in  $T$ , and the error incurred by approximating the integral in (1.5). For this reason, we would expect the order of the global error of the solver to be one less than the local error. However, the situation is more delicate because of the complicated manner in which the errors mix (see Figure 8.6). Our numerical results in Section 8 indicate that  $T$  converges with between  $O(h^2)$  and  $O(h^3)$  accuracy, and  $\nabla T$  converges with anywhere between  $O(h)$  and  $O(h^3)$  accuracy.

Regardless, we assume that we only know the values of the eikonal and some of its derivatives at the nodes  $\mathbf{x} \in \Omega_h$ . To obtain higher-order accuracy locally, we make use of piecewise Hermite elements. In particular, at each node  $\mathbf{x}$ , we approximate the *jet* of the eikonal; i.e.,  $\tau$  and a number of its derivatives [43]. If we compute the jet with sufficiently high accuracy when we set  $\text{state}(\mathbf{x}) \leftarrow \text{valid}$ , we will be in a position to approximate  $\tau$  using Hermite interpolation locally over  $\text{conv}(\mathbf{x}_1, \dots, \mathbf{x}_n)$ . We consider several variations on this idea.

**3.1. The general cost function.** Fix a point  $\hat{\mathbf{x}} \in \Omega_h$ , thinking of it as the *update point*. To compute  $T(\hat{\mathbf{x}})$ , we consider sets of **valid** nodes  $\{\mathbf{x}_1, \dots, \mathbf{x}_d\} \subseteq \text{nb}(\hat{\mathbf{x}})$ , where  $1 \leq d \leq n$ . The tuple of nodes  $(\hat{\mathbf{x}}, \mathbf{x}_1, \dots, \mathbf{x}_d)$  is an *update* of dimension  $d$ , and the collection of updates a *stencil*. We refer to the nodes  $\{\mathbf{x}_1, \dots, \mathbf{x}_d\}$  as the vertices of the *base of the update*. In some cases, such as on an unstructured mesh, stencils may vary with  $\hat{\mathbf{x}}$ . Sufficient conditions for the updates and stencils to be *monotone causal* have been studied [22]. In particular, the cone spanned by  $\{\mathbf{x}_1 - \hat{\mathbf{x}}, \dots, \mathbf{x}_n - \hat{\mathbf{x}}\}$  should fit inside the nonnegative orthant after being rotated [41, 42]. This is easily satisfied on a regular grid. For  $O(h)$  solvers that do not make use of gradient information, a variety of choices of stencils are monotone causal.

To describe a general update, without loss of generality we assume  $d = n$ , and assume that the update nodes are in general position. That is, if we choose  $n$  nodes from  $\{\hat{\mathbf{x}}, \mathbf{x}_1, \dots, \mathbf{x}_n\}$ , the remaining node does not lie in their convex hull. We assume that we have access to a sufficiently accurate approximation of  $\tau$  over  $\text{conv}(\mathbf{x}_1, \dots, \mathbf{x}_n)$ , call it  $\mathbb{T}$ . We distinguish between  $\mathbb{T}$  and  $T$  in the following way:  $\mathbb{T}$  denotes the local numerical approximation of  $\tau$  used by a particular update, while  $T$  denotes the global numerical approximation of  $\tau$ . The two may not be equal to each other. Indeed,  $T$  is in general only defined on  $\Omega_h$ , while  $\mathbb{T}$  is only defined on  $\text{conv}(\mathbf{x}_1, \dots, \mathbf{x}_n)$ .

Let  $\mathbf{x}_\lambda \in \text{conv}(\mathbf{x}_1, \dots, \mathbf{x}_n)$ , and let  $L = L_\lambda = \|\hat{\mathbf{x}} - \mathbf{x}_\lambda\|$ . Recall that  $\psi : [0, L] \rightarrow \Omega$  is the curve minimizing (1.5) for a particular choice of  $\mathbf{x}_\lambda$ . We approximate  $\psi$  with a cubic parametric curve  $\varphi : [0, L] \rightarrow \Omega$  such that:

$$(3.2) \quad \varphi(0) = \mathbf{x}_\lambda, \quad \varphi(L) = \hat{\mathbf{x}}, \quad \varphi'(0) \sim \mathbf{t}_\lambda, \quad \varphi'(L) \sim \hat{\mathbf{t}},$$

and where  $\mathbf{t}_\lambda$  and  $\hat{\mathbf{t}}$  are tangent vectors which enter as parameters. Note that  $\varphi'(0)$




FIG. 3.1. Two approaches to parametrizing a cubic curve approximating the characteristic  $\varphi$  leading from  $x_\lambda$  to  $\hat{x}$  when numerically minimizing Fermat's integral to compute  $T(\hat{x})$  and  $\nabla T(\hat{x})$ . Left:  $\varphi$  is a cubic parametric curve with boundary data set directly from  $x_\lambda$ ,  $\hat{x}$ ,  $t_\lambda$ , and  $\hat{t}$ . Right:  $\varphi$  is the graph of a function in the orthogonal complement of  $\text{range}(\ell')$ .

and  $\varphi'(L)$  may not be exactly equal to  $t_\lambda$  and  $\hat{t}$ .

We approximate the integral in (1.5) over  $\varphi$  using Simpson's rule. This gives the cost functional:

$$(3.3) \quad F(\varphi) = T(x_\lambda) + \frac{L}{6} \left[ s(x_\lambda) \|\varphi'(0)\| + 4s(\varphi_{1/2}) \|\varphi'_{1/2}\| + s(\hat{x}) \|\varphi'(L)\| \right],$$

where  $\varphi_{1/2} = \varphi(L/2)$  and  $\varphi'_{1/2} = \varphi'(L/2)$ . We have not yet made this well-defined. To do so, we must specify  $T$ ,  $t_\lambda$ , and  $\hat{t}$ . We describe several different ways of doing this in the following sections.

**3.2. Computing  $\nabla T(\hat{x})$ .** A minimizing extremal  $\psi$  of Fermat's integral is a characteristic of the eikonal equation. A simple but important consequence of this is that its tangent vector is locally parallel to  $\nabla \tau$ . Hence:

$$(3.4) \quad s(\psi(\sigma)) \frac{\psi'(\sigma)}{\|\psi'(\sigma)\|} = \nabla \tau(\psi(\sigma)).$$

After minimizing  $F$ , we will have found an optimal value of  $\hat{t}$ . We can then set:

$$(3.5) \quad \nabla T(\hat{x}) \leftarrow s(\hat{x}) \hat{t}.$$

This lets us march the gradient of the eikonal locally along with the eikonal itself.

**3.3. Parametrizing  $\varphi$ .** We consider two methods of choosing  $\varphi$  (see Figure 3.1). First, let  $\ell$  be interval connecting  $x_\lambda$  to  $\hat{x}$ , parametrized by arc length, and define:

$$(3.6) \quad \ell(\sigma) = x_\lambda + \sigma \ell', \quad \ell' = (\hat{x} - x_\lambda) / L_\lambda.$$

Using a cubic parametric curve. For one approach, we define:

$$(3.7) \quad \varphi(\sigma) = \ell(\sigma) + \delta\varphi(\sigma),$$

where  $\delta\varphi : [0, L] \rightarrow \Omega$  is a perturbation away from  $\ell$  that satisfies:

$$(3.8) \quad \delta\varphi(0) = 0, \quad \delta\varphi(L) = 0, \quad \delta\varphi'(0) = t_\lambda - \ell', \quad \delta\varphi'(L) = \hat{t} - \ell'.$$

We can explicitly write  $\delta\varphi$  as:

$$(3.9) \quad \delta\varphi(\sigma) = (\mathbf{t}_\lambda - \ell')K_0(\sigma) + (\hat{\mathbf{t}} - \ell')K_1(\sigma),$$

where  $K_0, K_1 : [0, L] \rightarrow \mathbb{R}$  are Hermite basis functions such that:

$$(3.10) \quad \begin{aligned} K_0(0) = 0 = K_0(L), & \quad K_1(0) = 0 = K_1(L), \\ K_0'(0) = 1 = K_0'(L), & \quad K_1'(0) = 0 = K_1'(L). \end{aligned}$$

Explicitly, these are given by:

$$(3.11) \quad K_0(\sigma) = \sigma - 2\frac{\sigma^2}{L} + \frac{\sigma^3}{L^2}, \quad K_1(\sigma) = \frac{-\sigma^2}{L} + \frac{\sigma^3}{L^2}.$$

Let  $\mathbf{t}_\lambda, \hat{\mathbf{t}} \in \mathbb{S}^{n-1}$  so that  $\|\mathbf{t}_\lambda\| = 1 = \|\hat{\mathbf{t}}\|$ . As  $L \rightarrow 0$ , this results in a curve that is approximately parametrized by arc length: i.e.,  $\|\varphi'(\sigma)\| \rightarrow 1$  for all  $\sigma$  such that  $0 \leq \sigma \leq L$  [16]. This simplifies the general cost function given by (3.3) to:

$$(3.12) \quad F(\varphi) = \mathsf{T}(\mathbf{x}_\lambda) + \frac{L}{6} \left[ s(\mathbf{x}_\lambda) + 4s(\varphi_{1/2})\|\varphi'_{1/2}\| + s(\hat{\mathbf{x}}) \right].$$

Using (3.11),  $\varphi_{1/2}$  and  $\varphi'_{1/2}$  can be written:

$$(3.13) \quad \varphi_{1/2} = \frac{\mathbf{x}_\lambda + \hat{\mathbf{x}}}{2} + \frac{L}{8}(\mathbf{t}_\lambda - \hat{\mathbf{t}}), \quad \varphi'_{1/2} = \frac{3}{2}\ell' - \frac{\mathbf{t}_\lambda + \hat{\mathbf{t}}}{4}.$$

Note that  $\varphi_{1/2} \sim (\mathbf{x}_\lambda + \hat{\mathbf{x}})/2$  and  $\varphi'_{1/2} \sim \ell'$  as  $L \rightarrow 0$  if we assume that the wavefront is well-approximated by a plane wave near the update, since in this case  $\mathbf{t}_\lambda \sim \hat{\mathbf{t}} \sim \ell'$ .

*Parametrizing  $\varphi$  as the graph of a function.* We can also define the perturbation away from  $\ell$  as the graph of a function; i.e., we assume that the perturbation is orthogonal to  $\ell'$ . Letting  $\mathbf{Q} \in \mathbb{R}^{n \times (n-1)}$  be an orthogonal matrix such that  $\mathbf{Q}^\top \ell' = 0$ , and letting  $\zeta : [0, L] \rightarrow \mathbb{R}^{n-1}$  be a curve specifying the components of the perturbation in this basis, we choose  $\delta\varphi(\sigma) = \mathbf{Q}\zeta(\sigma)$  so that:

$$(3.14) \quad \varphi(\sigma) = \ell(\sigma) + \mathbf{Q}\zeta(\sigma).$$

where  $\zeta(\sigma) = \mathbf{b}_0 K_0(\sigma) + \mathbf{b}_1 K_1(\sigma)$ . In this approach, instead of  $\hat{\mathbf{t}}$  and  $\mathbf{t}_\lambda$ , we optimize over  $\mathbf{b}_0, \mathbf{b}_1 \in \mathbb{R}^{n-1}$ . Now, noting that:

$$(3.15) \quad \|\varphi'(\sigma)\| = \sqrt{\|\ell'\|^2 + \|\mathbf{Q}\zeta'(\sigma)\|^2} = \sqrt{1 + \|\zeta'(\sigma)\|^2},$$

we can write the cost functional  $F$  as:

$$(3.16) \quad \begin{aligned} F(\varphi) = \mathsf{T}(\mathbf{x}_\lambda) + \frac{L}{6} \left[ s(\mathbf{x}_\lambda) \sqrt{1 + \|\mathbf{b}_0\|^2} \right. \\ \left. + 4s(\varphi_{1/2}) \sqrt{1 + \|(\mathbf{b}_0 + \mathbf{b}_1)/4\|^2} + s(\hat{\mathbf{x}}) \sqrt{1 + \|\mathbf{b}_1\|^2} \right]. \end{aligned}$$

*Trade-offs between the two parametrizations of  $\varphi$ .* When  $\varphi$  is a cubic parametric curve, we run into an interesting problem described in more detail by Floater [16]. In particular, the order of accuracy of  $\varphi$  in approximating  $\psi$  is limited by our parametrization of  $\varphi$ . If we parametrize  $\varphi$  over  $\sigma \in [0, 1]$  (that is *uniformly*), then the

interpolant is at most  $O(h^2)$  accurate. If we parametrize it using a *chordal parametrization*, i.e.  $\sigma \in [0, L]$ , then it is at most  $O(h^4)$  accurate. Indeed, any Hermite spline using a chordal parametrization over each of its segment is at most  $O(h^4)$  accurate globally. To design a higher order solver than this requires us to parametrize  $\varphi$  using a more accurate approximation of the arc length of  $\psi$  (consider, e.g., using a quintic parametric curve). On the other hand, if we parametrize  $\varphi$  as the graph of a function, we can directly apply Hermite interpolation theory [45], and there is no such obstacle.

**4. Different types of minimization problems.** In this section, we consider four different ways of using  $F$  to pose a minimization problem which would allow us to compute  $T(\hat{\mathbf{x}})$ . We note that each of these formulations is compatible with the version of  $F$  where we take  $\varphi$  to be a parametric curve *and* where we define it as the graph of a function orthogonal to  $\ell(\sigma)$ . Altogether, this leads to eight different JMMs.

**4.1. Determining  $t_\lambda$  by minimizing Fermat's integral.** Since the optimal  $\varphi$  is a characteristic of the eikonal equation, one approach to setting  $t_\lambda$  and  $\hat{\mathbf{t}}$  is to simply let them enter into the cost function as free parameters to be optimized over. This leads to the optimization problem:

$$(4.1) \quad \begin{aligned} & \text{minimize} && F(\mathbf{x}_\lambda, t_\lambda, \hat{\mathbf{t}}) \\ & \text{subject to} && \mathbf{x}_\lambda \in \text{conv}(\mathbf{x}_1, \dots, \mathbf{x}_n), \\ & && t_\lambda, \hat{\mathbf{t}} \in \mathbb{S}^{n-1}, \end{aligned}$$

if we parametrize  $\varphi$  as a curve; or, if we parametrize  $\varphi$  as the graph of a function:

$$(4.2) \quad \begin{aligned} & \text{minimize} && F(\mathbf{x}_\lambda, \mathbf{b}_0, \mathbf{b}_1) \\ & \text{subject to} && \mathbf{x}_\lambda \in \text{conv}(\mathbf{x}_1, \dots, \mathbf{x}_n), \\ & && \mathbf{b}_0, \mathbf{b}_1 \in \mathbb{R}^{n-1}, \end{aligned}$$

For a  $d$ -dimensional update, the domain of each of these minimization problems has dimension  $(d-1)(n-1)^2$ , since  $\dim(\text{conv}(\mathbf{x}_1, \dots, \mathbf{x}_d)) = d-1$ .

**4.2. Determining  $t_\lambda$  from the eikonal equation.** When we compute updates, we only require high-order accurate jets over  $\text{conv}(\mathbf{x}_1, \dots, \mathbf{x}_n)$ . This is a subset of  $\Omega$  of codimension one: an interval in 2D, or triangle in 3D. If we know  $T$  and  $\nabla T$  at the vertices of this set, then we can use Hermite interpolation to compute  $T$ . Unfortunately, this means that we can only approximate directional derivatives of  $T$  in the linear span of this set. To compute  $\nabla T$ , we need to recover the directional derivative normal to the facet.

Let  $\tilde{\mathbf{V}} \in \mathbb{R}^{n \times (n-1)}$  be an orthogonal matrix such that:

$$(4.3) \quad \text{range}(\tilde{\mathbf{V}}) = \text{range}([\mathbf{x}_2 - \mathbf{x}_1 \quad \dots \quad \mathbf{x}_n - \mathbf{x}_1]),$$

and let  $\mathbf{v} \in \mathbb{R}^n$  be a unit vector such that  $\tilde{\mathbf{V}}^\top \mathbf{v} = 0$ . Let  $\nabla_{\tilde{\mathbf{V}}}$  be the gradient restricted to the range of  $\tilde{\mathbf{V}}$ , and likewise let  $d_{\mathbf{v}}$  denote the  $\mathbf{v}$ -directional derivative. Then, from the eikonal equation, we have:

$$(4.4) \quad s(\mathbf{x})^2 = \|\nabla \tau(\mathbf{x})\|^2 = |d_{\mathbf{v}} \tau(\mathbf{x})|^2 + \|\nabla_{\tilde{\mathbf{V}}} \tau(\mathbf{x})\|^2.$$

To recover  $\nabla \tau(\mathbf{x})$ , first note that  $\nabla \tau(\mathbf{x})$  should point in the same direction as  $\ell'$ . Choosing  $\mathbf{v}$  so that  $\mathbf{v}^\top \ell' > 0$ , we get:

$$(4.5) \quad d_{\mathbf{v}} \tau(\mathbf{x}) = \sqrt{s(\mathbf{x})^2 - \|\nabla_{\tilde{\mathbf{V}}} \tau(\mathbf{x})\|^2}.$$

Letting  $\mathbf{V} = [\mathbf{v} \quad \tilde{\mathbf{V}}]$ , equation (4.5) combined with  $\nabla\tau(\mathbf{x}) = \mathbf{V}\nabla_{\mathbf{V}}\tau(\mathbf{x})$  gives us a means of recovering  $\nabla\tau(\mathbf{x})$  from  $\nabla_{\tilde{\mathbf{V}}}\tau(\mathbf{x})$ .

Using this technique, we can pose the following optimization problem:

$$(4.6) \quad \begin{aligned} & \text{minimize} && F(\mathbf{x}_\lambda, \hat{\mathbf{t}}) \\ & \text{subject to} && \mathbf{x}_\lambda \in \text{conv}(\mathbf{x}_1, \dots, \mathbf{x}_n), \\ & && \hat{\mathbf{t}} \in \mathbb{S}^{n-1}, \end{aligned}$$

or, optimizing over  $\mathbf{b}_1$  directly:

$$(4.7) \quad \begin{aligned} & \text{minimize} && F(\mathbf{x}_\lambda, \mathbf{b}_1) \\ & \text{subject to} && \mathbf{x}_\lambda \in \text{conv}(\mathbf{x}_1, \dots, \mathbf{x}_n), \\ & && \mathbf{b}_1 \in \mathbb{R}^{n-1} \end{aligned}$$

For each  $\mathbf{x}_\lambda$ , we set:

$$(4.8) \quad \mathbf{t}_\lambda \leftarrow \frac{\mathbf{V}\nabla_{\mathbf{V}}\Gamma(\mathbf{x}_\lambda)}{\|\mathbf{V}\nabla_{\mathbf{V}}\Gamma(\mathbf{x}_\lambda)\|}.$$

Note that the dimension of a  $d$ -dimensional update based on this minimization problem is  $(d-1)(n-1)$

**4.3. Determining  $\mathbf{t}_\lambda$  by marching cell-based interpolants.** Another approach is to march cells that approximate the jet of the eikonal at each point. For example, if we have constructed a finite element interpolant using valid data on a cell whose boundary contains  $\text{conv}(\mathbf{x}_1, \dots, \mathbf{x}_n)$  then we can evaluate its gradient to obtain:

$$(4.9) \quad \mathbf{t}_\lambda \leftarrow \frac{\nabla\Gamma(\mathbf{x}_\lambda)}{\|\nabla\Gamma(\mathbf{x}_\lambda)\|}.$$

We can combine this approach with the cost functional given by (4.6), albeit with a modified  $\mathbf{t}_\lambda$ . We elaborate on how we march cells in section 7. An advantage of this approach is that it allows one to simultaneously march the second partials of  $T$ .

**4.4. A simplified method using a quadratic curve.** In some cases, in particular if the speed of sound is linear, i.e.:

$$(4.10) \quad c(\mathbf{x}) = \frac{1}{s(\mathbf{x})} = c_0 + \mathbf{c}^\top \mathbf{x}, \quad c_0 \in \mathbb{R}, \quad \mathbf{c} \in \mathbb{R}^n,$$

the characteristic  $\psi$  is well-approximated by a quadratic. In this case, we again have a cost functional of the form (4.6).

If  $\varphi$  is parametrized as curve, we set  $\mathbf{t}_\lambda$  to be the reflection of  $\hat{\mathbf{t}}$  across  $\ell'$ :

$$(4.11) \quad \mathbf{t}_\lambda = -(\mathbf{I} - 2\ell'\ell'^\top)\hat{\mathbf{t}},$$

giving  $\mathbf{t}_\lambda + \hat{\mathbf{t}} = 2\ell'\ell'^\top\hat{\mathbf{t}}$  and  $\mathbf{t}_\lambda - \hat{\mathbf{t}} = -2(\mathbf{I} - \ell'\ell'^\top)\hat{\mathbf{t}}$ . Then:

$$(4.12) \quad \varphi_{1/2} = \frac{\mathbf{x}_\lambda + \hat{\mathbf{x}}}{2} - \frac{L}{4}(\mathbf{I} - \ell'\ell'^\top)\hat{\mathbf{t}}, \quad \varphi'_{1/2} = \frac{3 + \ell'^\top\hat{\mathbf{t}}}{2}\ell'.$$

This simplifies  $F$  given by (3.12) to:

$$(4.13) \quad F(\mathbf{x}_\lambda, \hat{\mathbf{t}}) = \Gamma(\mathbf{x}_\lambda) + \frac{L}{6} \left[ s(\mathbf{x}_\lambda) + 2(3 + \ell'^\top\hat{\mathbf{t}})s\left(\frac{\mathbf{x}_\lambda + \hat{\mathbf{x}}}{2} - \frac{L}{4}(\mathbf{I} - \ell'\ell'^\top)\hat{\mathbf{t}}\right) + s(\hat{\mathbf{x}}) \right].$$

If  $\varphi$  is parametrized as the graph of a function, then:

$$(4.14) \quad \zeta_{1/2} = \frac{\hat{\mathbf{x}} + \mathbf{x}_\lambda}{2} + \frac{L}{4} \mathbf{Q}^\top \hat{\mathbf{t}}, \quad \zeta' = 0,$$

simplifying the version of  $F$  in (3.16) to:

$$(4.15) \quad F(\mathbf{x}_\lambda, \hat{\mathbf{t}}) = \mathbb{T}(\mathbf{x}_\lambda) + \frac{L}{6} \left[ (s(\mathbf{x}_\lambda) + s(\hat{\mathbf{x}})) \sqrt{1 + \|\mathbf{b}_0\|^2} + 4s(\varphi_{1/2}) \right].$$

since  $\mathbf{Q}^\top \hat{\mathbf{t}} = \mathbf{b}_0 = -\mathbf{b}_1$ .

**4.5. Other approaches.** We tried two other approaches which failed to provide satisfactory results:

- A combination of the quadratic simplification in subsection 4.4 with the methods in subsections 4.2 or 4.3. In this case, we use our knowledge of  $\nabla T(\mathbf{x}_\lambda)$  along the base of the update to choose  $\hat{\mathbf{t}}$  and  $\mathbf{t}_\lambda$ . This reduces the dimensionality of the cost function to  $d - 1$ . However, except for in special cases (e.g.  $s \equiv 1$ ), this propagates errors in a manner that causes the solver to diverge; or, at best, allows it converge with  $O(h)$  accuracy. We note that if  $s \equiv 1$ , still simpler methods can be used, so this combination of approaches does not seem to be useful.
- We can extract not only  $\mathbf{t}_\lambda$  from the Hermite interpolant on  $\text{conv}(\mathbf{x}_1, \dots, \mathbf{x}_n)$ , but also  $\varphi''(0)$ . From the Euler-Lagrange equations for the eikonal equation, we obtain:

$$(4.16) \quad \mathbf{Q}^\top \nabla s(\mathbf{x}_\lambda) = s(\mathbf{x}_\lambda) \frac{\mathbf{Q}^\top \varphi''(0)}{\|\varphi'(0)\|}.$$

With  $\varphi$  parametrized as a graph, we have  $\mathbf{Q}^\top \varphi''(0) = \zeta''(0)$ , giving:

$$(4.17) \quad \zeta''(0) = \frac{\mathbf{Q}^\top \nabla s(\mathbf{x}_\lambda) (1 + \|\mathbf{Q}^\top \mathbf{x}_\lambda\|^2)}{s(\mathbf{x}_\lambda)}.$$

This completely defines  $\varphi$  as the graph of a cubic polynomial using the graph parametrization. Unfortunately, this method diverges for the same reason as the method described in the previous bullet.

**4.6. A warm start.** Certain of the optimization problems in the preceding section are clearly nonconvex; e.g., (4.1) is nonconvex since its domain, the product of  $\text{conv}(\mathbf{x}_1, \dots, \mathbf{x}_n)$  and two copies of  $\mathbb{S}^{n-1}$ , is nonconvex. For  $h$  small, if  $\mathbf{x}_\lambda$  is nearly optimal, then the optimal local ray  $\varphi$  should be close to  $\ell$ , the straight line segment connecting  $\mathbf{x}_\lambda$  and  $\hat{\mathbf{x}}$ . This suggests an approach to finding an initial iterate (a warm start) for (4.1) or the other optimization problems considered (i.e., (4.2), (4.6), and (4.7)). In 2D, following a similar approach to the simplified midpoint rule (denoted “mp0”) rule used in our earlier work on ordered line integral methods (OLIMs) for the solving the eikonal equation [30], we let  $\mathbb{T}$  be a cubic Hermite polynomial approximation of  $\tau$  over  $\lambda \in [0, 1]$  and approximate (1.5) by solving:

$$(4.18) \quad \begin{aligned} & \text{minimize} && \mathbb{T}(\lambda) + \frac{L}{2} (s(\mathbf{x}_\lambda) + s(\hat{\mathbf{x}})), \\ & \text{subject to} && 0 \leq \lambda \leq 1. \end{aligned}$$

After solving this problem, we can compute an initial iterate for (4.1) from  $\lambda^*$ , the optimum of (4.18). For example, we set  $\mathbf{x}_\lambda \leftarrow \mathbf{x}_{\lambda^*}$ ; we set  $\mathbf{t}_\lambda$  using  $\lambda^*$  and one of the approaches outlined in the preceding sections; and, if required, we set  $\hat{\mathbf{t}} \leftarrow (\hat{\mathbf{x}} - \mathbf{x}_{\lambda^*}) / \|\hat{\mathbf{x}} - \mathbf{x}_{\lambda^*}\|$ . In practice, (4.18) can be solved rapidly and robustly using a rootfinder.

**4.7. Optimization algorithms.** We do not dwell on the details of how to numerically solve the minimization problems in the preceding sections. We make some general observations:

- These optimization problems are very easy to solve—what’s costly is that we have to solve  $O(N)$  of them. As  $h \rightarrow 0$ , they are strictly convex and well-behaved. Empirically, Newton’s method converges in  $O(1)$  steps (typically fewer than 5 with a well-chosen warm start—see section 4.6). We leave a detailed comparison of different approaches to numerically solving these optimization problems for future work.
- The gradients and Hessians of these cost functions are somewhat complicated. Programming them can be tricky and tedious, suggesting that automatic differentiation may be a worthwhile approach [26, 19].
- The constraint  $\mathbf{x}_\lambda \in \text{conv}(\mathbf{x}_1, \dots, \mathbf{x}_n)$  corresponds to a set of linear inequality constraints, which are simple to incorporate. Because of the form of these constraints, checking the KKT conditions at the boundary is cheap and easy [30, 49]. See the next section on skipping updates.
- The constraints  $\mathbf{t}_\lambda, \hat{\mathbf{t}} \in \mathbb{S}^{n-1}$  are nonlinear; however, they can be eliminated. If  $n = 2$ , then we can set  $\hat{\mathbf{t}} = (\cos(\theta), \sin(\theta))$ , letting  $\theta \in \mathbb{R}$ . For  $n > 2$ , we can use a Riemannian Newton’s method for minimization on  $\mathbb{S}^{n-1}$ , which is simple to implement and known to converge superlinearly [1]. Alternatively, we could use spherical coordinates, although the expressions become unwieldy.

**5. Hierarchical update algorithms.** Away from shocks, where multiple wavefronts collide, exactly one characteristic will pass through a point  $\hat{\mathbf{x}}$ . When we minimize  $F$  over each update in the stencil, the characteristic will pass through the base of the minimizing update, or possibly through the boundary of several adjacent updates. We can use this fact to sequence the updates that are performed to design a work-efficient solver. In our previous work on *ordered line integral methods (OLIMs)*, we explored variations of this idea [30, 49]. An approach that works well is the *bottom-up* update algorithm.

To fix the idea in 3D, consider  $\text{nb}(\mathbf{x})$  as shown in Figure 5.1, for which  $|\text{nb}(\mathbf{x})| = 26$ . There are 26 “line” updates, where  $d = 1$ . To start with, each `valid` line update is done, and  $\mathbf{x}_1$  for the minimizing line update is recorded. Next, we fix  $\mathbf{x}_1$  and perform “triangle” updates ( $d = 2$ ) where  $\mathbf{x}_2$  is varying. In this case, we can restrict the number of triangle updates that are done by assuming either that  $(\mathbf{x}_1, \mathbf{x}_2)$  is an edge of mesh discretizing the surface of the 3D stencil shown in Figure 5.1, or that  $\|\mathbf{x}_1 - \mathbf{x}_2\|$  is small enough (measuring the distance of these two points in different norms leads to a different number of triangle updates—we find the  $\ell_1$  norm to work well). Finally, we fix  $\mathbf{x}_2$  corresponding to the minimizing triangle update, and do tetrahedron updates containing  $\mathbf{x}_1$  and  $\mathbf{x}_2$ . Throughout this process,  $\mathbf{x}_1, \mathbf{x}_2$ , and  $\mathbf{x}_3$  must all be `valid`.

We emphasize that our work-efficient OLIM update algorithms work equally well for the class of algorithms developed here. The main differences between the JMMs studied here and the earlier OLIMs are the cost functionals and the way approximate  $T$ .




FIG. 5.1. The neighborhoods typically used by semi-Lagrangian solvers in 2D and 3D on a regular grid. These are the stencils used by Tsitsiklis’s algorithm and two of the OLIM stencils [46, 30]. Left: “olim8” in  $\mathbb{R}^2$ . This is the 8-point stencil used in this paper. Right: “olim26” in  $\mathbb{R}^3$ .

**6. Initialization methods.** A common problem with the convergence of numerical methods for solving the eikonal equation concerns how to treat rarefaction fans. Our numerical tests consist of point source problems, around which a rarefaction forms. A standard approach is to introduce the factored eikonal equation [17, 24, 31]. If a point source is located at  $\mathbf{x}^\circ \in \Omega_h$  and if we set  $\Gamma_h = \{\mathbf{x}^\circ\}$ , then we let  $d(\mathbf{x}) = \|\mathbf{x} - \mathbf{x}^\circ\|$  and use the ansatz:

$$(6.1) \quad \tau(\mathbf{x}) = z(\mathbf{x}) + d(\mathbf{x}), \quad \mathbf{x} \in \Omega.$$

We insert this into the eikonal equation, modifying our numerical methods as necessary, and solve for  $z(\mathbf{x})$  instead. This is not complicated—see our previous work on OLIMs for solving the eikonal equation to see how the cost functions should generally be modified [30, 49].

Yet another approach would be to solve the characteristic equations to high-order for each  $\mathbf{x}$  in such a ball. This would require solving  $O(N)$  boundary value problems, each discretized into  $O(N^{1/3})$  intervals, resulting in an  $O(N^{4/3})$  cost overall (albeit with a very small constant). One issue with this approach is that it only works well if the ball surrounding  $\mathbf{x}^\circ$  is contained in the interior of  $\Omega$ . For our numerical experiments, we simply initialize  $T$  and  $\nabla T$  to the correct, ground truth values in a ball or box of constant size centered at  $\mathbf{x}^\circ$ .

**7. Cell marching.** Of particular interest is solving the transport equation governing the amplitude  $\alpha$  while simultaneously solving the eikonal equation. Equation (1.4) can be solved using upwind finite differences [5] or paraxial raytracing [29]. We prefer the latter approach since it can be done locally, using the characteristic path  $\varphi$  recovered when computing  $T(\hat{\mathbf{x}})$ . Either approach requires accurate second derivative information (we need  $\Delta T$  for upwind finite differences, or  $\nabla^2 T$  for paraxial raytracing).

For the purposes of explanation and our numerical tests, we consider a rectilinear grid with square cells in  $\mathbb{R}^2$ . On each cell, our goal is to build a bicubic interpolant, approximating  $T(\mathbf{x})$ . This requires knowing  $T$ ,  $\nabla T$ , and  $T_{xy}$  at each cell corner. If we know these values with  $O(h^{4-p})$  accuracy, where  $p$  is the order of the derivative, then




FIG. 7.1. Cell-based interpolation. To approximate the mixed second partials of a function with  $O(h^2)$  accuracy from  $O(h^3)$  accurate gradient values available at the corners of a cell, the following method of using central differences to approximate the mixed partials at the midpoints of the edges of the cell, followed by bilinear extrapolation, can be used.




FIG. 7.2. Local cell marching. After computing values of  $T_{xy}$  as shown in Figure 7.1 (left), to ensure continuity of the global interpolant, nodal values incident on the newly valid cell (containing  $\mathbf{x}_0$ ) can be recomputed by averaging over  $T_{xy}$  values taken from incident valid cells (middle). Finally, a bicubic cell-based interpolant is constructed (right).

the bicubic is  $O(h^{4-p})$  accurate over the cell. So far, we have described an algorithm that marches  $T$  and  $\nabla T$ , which together constitute the *total 1-jet*. We now show how  $T_{xy}$  can also be marched, allowing us to march the *partial 1-jet*.<sup>1</sup>

Let  $\mathbf{x}_{ij}$  with  $(i, j) \in \{0, 1\}^2$  denote the corners of a square cell with sides of length  $h$ , and assume that we know  $\nabla T(\mathbf{x}_{ij})$  with  $O(h^3)$  accuracy. We can use the following approach to estimate  $T_{xy}(\mathbf{x}_{ij})$  at each corner:

- First, at the midpoints of the edges oriented in the  $x$  direction (resp.,  $y$  direction), approximate  $T_{xy}$  using the central differences involving  $T_y$  (resp.,  $T_x$ ) at the endpoints. This approximation is  $O(h^2)$  accurate at the midpoints.
- Use bilinear extrapolation to reevaluate  $T_{xy}$  at the corners of the cell, yielding  $T_{xy}(\mathbf{x}_{ij})$ , also with  $O(h^2)$  accuracy.

This procedure is illustrated in Figure 7.1.

One issue with this approach is that it results in a piecewise interpolant that is only  $C^1$  globally. That is, if we estimate the value of  $T_{xy}$  at a corner from each of the cells which are incident upon it, we will get different values in general. To compute a globally  $C^2$  piecewise interpolant, we can average  $T_{xy}$  values over incident valid cells, where we define a **valid** cell to be a cell whose vertices are all **valid**. How to

<sup>1</sup>The total  $k$ -jet of a function  $f$  is the set  $\{\partial^\alpha f\}_\alpha$ , where  $\|\alpha\|_1 \leq k$ ; the partial  $k$ -jet is  $\{\partial^\alpha f\}_\alpha$  where  $\|\alpha\|_\infty \leq k$ .

do this is shown in Figure 7.2.

The idea of approximating the partial 1-jet from the total 1-jet in an optimally local fashion by combining central differences with bilinear extrapolation, and averaging nodal values over adjacent cells to increase the degree of continuity of the interpolant, is borrowed from Seibold et al. [37]. However, applying this idea in this context, and doing the averaging in an upwind fashion is novel.

The scheme arrived at in this way is no longer optimally local. However, the sequence of operations described here can be done on an unstructured triangle or tetrahedron mesh. This makes this approach suitable for use with an unstructured mesh that conforms to a complicated boundary. We should mention here that our approach to estimating  $T_{xy}$  is referred to as *twist estimation* in the *computer-aided design* (CAD) community [15], where other approaches have been proposed [10, 20]. We leave adapting these ideas to the present context for future work.

**7.1. Marching the amplitude.** In this section we show how to compute a numerical approximation of  $\alpha$ , denoted  $A : \Omega_h \rightarrow \mathbb{R}$ . One simple approach would be to discretize (1.4) using upwind finite differences and compute  $A(\hat{\mathbf{x}})$  using valid nodes after  $T(\hat{\mathbf{x}})$  and  $\nabla T(\hat{\mathbf{x}})$ . One potential shortcoming of this approach is that  $A$  is singular at caustics. Instead, we will explore using paraxial raytracing to compute  $A$  in this section [28]. The background material on paraxial raytracing used in this section can be found in more detail in M. Popov’s book [29].

The basic idea of paraxial raytracing is to consider a fixed, central ray, which we denote  $\varphi_0$ , and a surrounding tube of rays, parametrized by:

$$(7.1) \quad \varphi(\sigma, \mathbf{q}) = \varphi_0(\sigma) + \mathbf{E}(\sigma)\mathbf{q},$$

where  $\mathbf{E} : [0, L] \rightarrow \mathbb{R}^{n \times (n-1)}$  is an orthogonal matrix such that  $\mathbf{E}^\top \varphi_0' \equiv 0$ . For each  $\mathbf{q}$ , the corresponding ray should satisfy the Euler-Lagrange equations for (1.1). If we let  $c_0(\sigma) = c(\varphi_0(\sigma))$ , where  $c = 1/s$ , then  $\mathbf{q}$  along with the conjugate momenta  $\mathbf{p}$  (the exact form of which is not important in this instance) will satisfy:

$$(7.2) \quad \begin{bmatrix} d\mathbf{q}/d\sigma \\ d\mathbf{p}/d\sigma \end{bmatrix} = \begin{bmatrix} c_0(\sigma)\mathbf{I} \\ \frac{-1}{c_0(\sigma)^2} \frac{\partial^2 c}{\partial \mathbf{q} \partial \mathbf{q}^\top} \Big|_{\mathbf{q}=0} \end{bmatrix} \begin{bmatrix} \mathbf{q} \\ \mathbf{p} \end{bmatrix}.$$

If we let  $\mathbf{Q}(\sigma), \mathbf{P}(\sigma) : [0, L] \rightarrow \mathbb{R}^{(n-1) \times (n-1)}$  be a linearly independent set of solutions to (7.2), then along the central ray, the amplitude satisfies:

$$(7.3) \quad A(\varphi_0(\sigma)) = \sqrt{\frac{c_0(\sigma)}{|\det(\mathbf{Q}(\sigma))|}} A(\varphi_0(0)).$$

Note that when we compute an update, we obtain a cubic path  $\varphi$  approximating a ray of (1.1), such that  $\varphi(0) = \mathbf{x}_\lambda$  and  $\varphi(L) = \hat{\mathbf{x}}$ .

The quantity  $|\det(\mathbf{Q}(\sigma))|$  is known as the *geometric spreading* along the ray tube. We denote it  $J(\sigma)$ . Letting  $A$  denote a polynomial approximation of  $A$  off of the grid nodes in  $\Omega_h$ , using (7.3), we can compute  $A(\hat{\mathbf{x}})$  from:

$$(7.4) \quad A(\hat{\mathbf{x}}) = \sqrt{\frac{c_0(L_\lambda)}{|\det(\mathbf{Q}(L_\lambda))|}} A(\mathbf{x}_\lambda) = \sqrt{\frac{c(\hat{\mathbf{x}})}{J(\mathbf{x}_\lambda)}} A(\mathbf{x}_\lambda).$$

Since this depends on  $\mathbf{Q}(L)$ , we must solve (7.2) along  $\varphi$ , requiring us to provide initial conditions at  $\sigma = 0$ . Note that if we set  $\sigma = 0$  in (7.3), we can see that

$|\det(\mathbf{Q}(0))| = c_0(0)$  is necessary. A simple choice for the initial conditions for  $\mathbf{Q}$  is  $\mathbf{Q}(0) = c_0(0)^{1/n} \mathbf{I}$ . This assumes that we aren't too close to a point source, where  $A$  is singular.

To find initial conditions for  $\mathbf{P}$ , first expand  $\tau$  in a Taylor series orthogonal to the central ray, i.e. in the coordinates  $\mathbf{q}$ . Doing this, we find that:

$$(7.5) \quad \tau(\varphi(\sigma, \mathbf{q})) = \tau(\varphi_0(\sigma)) + \frac{1}{2} \mathbf{q}^\top \left. \frac{\partial^2 \tau}{\partial \mathbf{q} \partial \mathbf{q}^\top} \right|_{\mathbf{q}=0} \mathbf{q} + O(\mathbf{q}^3).$$

In this Taylor expansion, the linear term disappears since the rays and wavefronts are orthogonal. If we let:

$$(7.6) \quad \mathbf{\Gamma} = \left. \frac{\partial^2 \tau}{\partial \mathbf{q} \partial \mathbf{q}^\top} \right|_{\mathbf{q}=0},$$

we find that  $\mathbf{\Gamma}$  satisfies the matrix Ricatti equation:

$$(7.7) \quad \frac{d\mathbf{\Gamma}}{d\sigma} + c_0 \mathbf{\Gamma}^2 + \frac{1}{c_0^2} \left. \frac{\partial^2 c}{\partial \mathbf{q} \partial \mathbf{q}^\top} \right|_{\mathbf{q}=0} = 0.$$

The standard way to solve (7.7) is to use the ansatz  $\mathbf{\Gamma} = \mathbf{P}\mathbf{Q}^{-1}$ , which, indeed, leads us back to (7.2). However, this viewpoint furnishes us with the initial conditions for  $\mathbf{P}$ , since  $\mathbf{\Gamma}(0)$  can now be readily computed from  $\nabla^2 \mathbf{T}(\mathbf{x}_\lambda)$ .

*Marching the amplitude of a linear speed of sound.* As a simple but important test case, we consider a problem with a constant speed of sound, i.e.:

$$(7.8) \quad s(\mathbf{x}) = \frac{1}{c(\mathbf{x})}, \quad c(\mathbf{x}) = v_0 + \mathbf{v}^\top \mathbf{x}.$$

In this case, (7.2) simplifies considerably since  $\nabla^2 c \equiv 0$ , implying  $\mathbf{P}(\sigma) \equiv \mathbf{P}(0) = \mathbf{\Gamma}(0)\mathbf{Q}(0)$ . From this, we can integrate  $d\mathbf{Q}/d\sigma$  from 0 to  $L$  to obtain:

$$(7.9) \quad \mathbf{Q}(L) = \left[ \mathbf{I} + \left( \int_0^L c(\varphi(\sigma)) d\sigma \right) \mathbf{\Gamma}(0) \right] \mathbf{Q}(0).$$

Denote the integral in this expression for  $\mathbf{Q}(L)$  by  $\epsilon$ . To evaluate  $\epsilon$  approximately, we can apply the trapezoid rule to get:

$$(7.10) \quad \epsilon = \int_0^L c(\varphi(\sigma)) d\sigma = L \cdot (v_0 + \mathbf{v}^\top (\hat{\mathbf{x}} + \mathbf{x}_\lambda)/2) + O(L^2),$$

which implies that  $|\epsilon| = O(L)$ . The fact that the error is  $O(L^2)$  in this case follows from usual error bound for the trapezoid rule and the fact that  $\max_{0 \leq \sigma \leq L} |\varphi''(\sigma)| = O(L^{-1})$ , by our choice of parametrization.

We would like to develop a simple update rule for the geometric spreading. First, note that the determinant satisfies the following identity:

$$(7.11) \quad \det(\mathbf{I} + \epsilon \mathbf{\Gamma}(0)) = 1 + \epsilon \operatorname{tr}(\mathbf{\Gamma}(0)) + O(\epsilon^2).$$

Next, recall that  $\mathbf{E}(0)$  is an orthogonal matrix such that  $\mathbf{E}(0)^\top \varphi'_0 = \mathbf{E}(0)^\top \mathbf{t}_\lambda = 0$ . Let  $\mathbf{U} = [\mathbf{E}(0) \quad \mathbf{t}_\lambda]$  and write:

$$(7.12) \quad \operatorname{tr} \nabla^2 T(\mathbf{x}_\lambda) = \operatorname{tr} \mathbf{U}^\top \nabla^2 T(\mathbf{x}_\lambda) \mathbf{U} = \operatorname{tr} \mathbf{E}(0)^\top \nabla^2 T(\mathbf{x}_\lambda) \mathbf{E}(0) + \mathbf{t}_\lambda^\top \nabla^2 T(\mathbf{x}_\lambda) \mathbf{t}_\lambda.$$

By definition,  $\mathbf{\Gamma}(0) = \mathbf{E}(0)^\top \nabla^2 T(\mathbf{x}_\lambda) \mathbf{E}(0)$ . Taking the gradient of (1.1), we get:

$$(7.13) \quad \nabla^2 T(\mathbf{x}_\lambda) \nabla T(\mathbf{x}_\lambda) = s(\mathbf{x}_\lambda) \nabla s(\mathbf{x}_\lambda),$$

which leads immediately to:

$$(7.14) \quad \mathbf{t}_\lambda^\top \nabla^2 T(\mathbf{x}_\lambda) \mathbf{t}_\lambda = \mathbf{t}_\lambda^\top \nabla s(\mathbf{x}_\lambda),$$

noting that  $\mathbf{t}_\lambda = \nabla T(\mathbf{x}_\lambda) / \|\nabla T(\mathbf{x}_\lambda)\|$ . Combining (7.12) and (7.14) gives:

$$(7.15) \quad \text{tr } \mathbf{\Gamma}(0) = \Delta T(\mathbf{x}_\lambda) - \mathbf{t}_\lambda^\top \nabla s(\mathbf{x}_\lambda),$$

since  $\text{tr } \nabla^2 T(\mathbf{x}_\lambda) = \Delta T(\mathbf{x}_\lambda)$ . This gives the following update for  $J$ :

$$(7.16) \quad J(\hat{\mathbf{x}}) = \left| 1 + \epsilon \cdot (\Delta T(\mathbf{x}_\lambda) - \mathbf{t}_\lambda^\top \nabla s(\mathbf{x}_\lambda)) \right| \cdot \mathbf{J}(\mathbf{x}_\lambda).$$

Here,  $\mathbf{J}$  denotes a local polynomial approximation to  $J$ . This can be computed directly from data immediately available after solving the optimization problem that determines  $T(\hat{\mathbf{x}})$  and  $\nabla T(\hat{\mathbf{x}})$ .

*Initial data for  $J$  and  $A$ .* Determining the initial data for the amplitude is involved [3, 4, 29, 32], and detailed consideration of this problem is outside the scope of this work. Instead, we note that for a point source in 2D, the following hold approximately near the point source:

$$(7.17) \quad J(\mathbf{x}) \sim |\mathbf{x}|, \quad A(\mathbf{x}) \sim \frac{e^{i\pi/4}}{2\sqrt{2\pi\omega}} \sqrt{\frac{c(\mathbf{x})}{J(\mathbf{x})}}.$$

See Popov for quick derivations of these approximations [29]. In our test problems, we initialize  $J$  to  $|\mathbf{x}|$  near the point source, march  $J$  according to (7.16) where  $\mathbf{J} = (1 - \lambda)J(\mathbf{x}_1) + \lambda J(\mathbf{x}_2)$ , and compute the final amplitude from:

$$(7.18) \quad A(\mathbf{x}) = \frac{e^{i\pi/4}}{2\sqrt{2\pi\omega}} \sqrt{\frac{c(\mathbf{x})}{J(\mathbf{x})}}.$$

We emphasize that this is only valid for two-dimensional problems. The same sort of approach can be used for 3D problems, but (7.17) must be modified.

*Marching the amplitude for more general slowness functions.* The update given by (7.16) is valid if we approximate the speed function  $c = 1/s$  with a piecewise linear function with nodal values taken from  $c(\mathbf{x})$ , where  $\mathbf{x} \in \Omega_h$ . This should be a reasonable thing to do, since the update rule given by (7.16) in this case appears to be  $O(h^2)$  accurate. Since the accuracy of  $\nabla^2 T$  computed by our method is limited, we should not expect to be able to obtain much better than  $O(h)$  accuracy for  $J$ . That said, a more accurate update for  $J$  could be obtained by numerically integrating (7.2).

**8. Numerical experiments.** In this section, we first present a variety of test problems which differ primarily in the choice of slowness function  $s$ . The choices of  $s$  range from simple, such as  $s \equiv 1$  (an overly simplified but reasonable choice for speed of sound in room acoustics), to more strongly varying. We then present experimental results for our different JMMs as applied to these different slowness functions, demonstrating the significant effect the choice of  $s$  has on solver accuracy. The solvers used in these experiments are:

- JMM1:  $\varphi$  is approximated using a cubic curve, and tangent vectors are found by solving 4.1.
- JMM2:  $\varphi$  is approximated using a cubic curve, with  $\hat{\mathbf{t}}$  optimized from 4.6 and  $\mathbf{t}_\lambda$  found from Hermite interpolation at the base of the update.
- JMM3:  $\varphi$  is approximated using a quadratic curve, with its tangent vectors being found by optimizing.
- JMM4: JMM2 combined with the cell-marching method described in section 7.

We also plot the same results obtained by the FMM [38] and `olim8_mp0` [30]. We do not include least squares fits for these solvers in our tables. They are mostly  $O(h)$ , with some exceptions for  $\nabla T$  as computed by the FMM.

We note that  $s$  does not significantly affect the runtime of any of our solvers—formally, our solvers run in  $O(|\Omega_h| \log |\Omega_h|)$  time, where the constant factors are essentially insensitive to the choice of  $s$ . We note that the cost of updating the heap is very small compared to the cost of doing updates. Since only  $|\Omega_h|$  updates must be computed, the CPU time of the solver effectively scales like  $O(|\Omega_h|)$  for all problem sizes considered in this paper.

**8.1. Test problems.** In this section, we provide details for the test problems used in our numerical tests.

*Constant slowness with a point source.* For this problem, the slowness and solution are given by:

$$(8.1) \quad s \equiv 1, \quad \tau(\mathbf{x}) = \|\mathbf{x}\|.$$

We take the domain to be  $\Omega = [-1, 1] \times [-1, 1] \subseteq \mathbb{R}^2$ . To control the size of the discretized domain, we let  $M > 0$  be an integer and set  $h = 1/M$ , from which we define  $\Omega_h$  accordingly. We place a point source at  $\mathbf{x}^\circ = (0, 0) \in \Omega_h$ . The set of initial boundary data locations given by is  $\Gamma_h = \{\mathbf{x}^\circ\}$ , with boundary conditions given by  $g(\mathbf{x}^\circ) = 0$ .

*Linear speed with a point source (#1).* Our next test problem has a linear velocity profile. This might model the variation in the speed of sound due to a linear temperature gradient (e.g., in a large room). The slowness is given by [17, 44]:

$$(8.2) \quad s(\mathbf{x}) = \left[ \frac{1}{s_0} + \mathbf{v}^\top \mathbf{x} \right]^{-1},$$

where  $s_0 > 0$ , and  $\mathbf{v} \in \mathbb{R}^2$  are parameters. The solution is given by:

$$(8.3) \quad \tau(\mathbf{x}) = \frac{1}{\|\mathbf{v}\|} \cosh^{-1} \left( 1 + \frac{1}{2} s_0 s(\mathbf{x}) \|\mathbf{v}\|^2 \|\mathbf{x}\|^2 \right).$$

For our first test with a linear speed function, we take  $s_0 = 1$  and  $\mathbf{v} = (0.133, -0.0933)$ . For this problem,  $\Omega = [-1, 1] \times [-1, 1]$ ,  $\Gamma_h = \{\mathbf{x}^\circ\}$ , and  $g(\mathbf{x}^\circ) = 0$ .

*Linear speed with a point source (#2).* For our second linear speed test problem, we set  $s_0 = 2$  and  $\mathbf{v} = (0.5, 0)$  as in [31]. For this problem, we let  $\Omega = [0, 1] \times [0, 1]$ , discretize into  $M$  nodes along each axis, and define  $\Omega_h$  accordingly (i.e.,  $|\Omega_h| = M^2$ , with  $M = h^{-1}$ ). We take  $\mathbf{x}^\circ$ ,  $\Gamma_h$ , and  $g$  to be same as in the previous two test problems.

*A nonlinear slowness function involving a sine function.* For  $\mathbf{x} = (x_1, x_2)$ , we set:

$$(8.4) \quad \tau(\mathbf{x}) = \frac{x_1^2}{2} + 2 \sin \left( \frac{x_1 + x_2}{2} \right)^2.$$




FIG. 8.1. A test problem with two point sources. The domain is  $\Omega = [0, 1]^2$  discretized into a  $33 \times 33$  grid. The top row contains computed values and the bottom row contains signed errors. We overlay several contours of the numerically computed eikonal,  $T$ , for context. The two point sources are separated by a shockline, which is accurately localized even for this extremely coarse mesh. E.g., if we compared values of  $T_x$  or  $T_y$  on the “wrong” side of the shockline, we would observe  $O(1)$  error. Left:  $T$ . Middle:  $T_x$ . Right:  $T_y$ .

This eikonal has a unique minimum,  $\tau(0, 0) = 0$ , and is strictly convex in  $\Omega = [-1, 1] \times [-1, 1]$ . This lets us determine the slowness from the eikonal equation, giving:

$$(8.5) \quad s(\mathbf{x}) = \sqrt{\sin(x_1 + x_2)^2 + (x_1 + \sin(x_1 + x_2))^2}.$$

For this test problem, we take  $\Gamma_h$  and  $\Omega_h$  as in the constant slowness point source problem.

*Sloth.* A slowness function called “sloth” (jargon from geophysics) is taken from Example 1 of Fomel et al. [17]:

$$(8.6) \quad s(\mathbf{x}) = \sqrt{s_0^2 + 2\mathbf{v}^\top \mathbf{x}}.$$

For our test with this slowness function, we set  $s_0 = 2$ , and  $\mathbf{v} = (0, -3)$ . In this case, to avoid shadow zones formed by caustics, we take  $\Omega = [0, \frac{1}{2}] \times [0, \frac{1}{2}]$ . The discretized domain and boundary data are determined analogously to the earlier cases.

*Two point sources.* We consider a linear speed function:

$$(8.7) \quad c(\mathbf{x}) = \frac{1}{s(\mathbf{x})} = 2 + 5x_1 + 20x_2,$$

with point sources located at  $\mathbf{x}_0 = (0, 0)$  and  $\mathbf{x}_1 = (0.8, 0)$ . This is the same problem considered in Figure 4 of Qi and Vladimirovsky [31]. We use (8.3) to compute the groundtruth eikonal for each point source. If we let  $\tau_0$  and  $\tau_1$  denote the eikonals for each point source problem considered individually, then the combined eikonal is:

$$(8.8) \quad \tau(\mathbf{x}) = \min(\tau_0(\mathbf{x}), \tau_1(\mathbf{x})).$$

We can easily see that all derivatives of  $\tau$  are undefined on the shockline, where  $\tau_0 \equiv \tau_1$ . However, away from the shockline, the derivatives are well-defined and




FIG. 8.2. A single reflection from a wall in a semi-infinite domain for  $\omega = 1000$  (we consider only the top edge of the boundary to be reflecting). Left:  $T$  for the incident and reflected fields. The reflected field satisfies a specular reflection condition along the edge of the domain. Right: the real part of the approximation to the solution of the Helmholtz equation so obtained, given by (8.15).

smooth: if  $j = \arg \min_i \tau_i(\mathbf{x})$ , then  $D\tau \equiv D\tau_j$ . We solve this problem on  $\Omega = [0, 1]^2$ , so that  $\Omega_h$  is a grid with  $N = 33$  nodes in each direction. The results are shown in Figure 8.1.

For this problem, we note that the shockline is localized sufficiently well: all nodes  $\mathbf{x} \in \Omega_h$  are on the correct side, so that when we compute the errors, each nodal value is compared with the correct branch of  $\tau = \min(\tau_0, \tau_1)$ . Note the use of an extremely coarse mesh. The mesh used here is coarser than any mesh used in Qi and Vladimirov's Figure 4, but the eikonal achieves a smaller maximum error than all of their test problems.

A *single reflection*. We additionally include a simple test for computing multiple arrivals. For the linear speed function:

$$(8.9) \quad c(\mathbf{x}) = \frac{1}{s(\mathbf{x})} = 2 + 5x_1 + 7x_2,$$

we solve (1.1) on  $\Omega = [0, 1]^2$ , discretized into  $N = 101$  nodes in each direction. We place a point source at  $\mathbf{x} = (0, 0)$  and compute the eikonal, which we denote  $\tau_{\text{in}}$ . We then restrict  $\tau_{\text{in}}$  and  $\nabla\tau_{\text{in}}$  (after reflection) to the set  $\Gamma = [0, 1] \times \{1\}$  (the top edge of the domain), and solve the reflected eikonal equation:

$$(8.10) \quad \|\nabla\tau_{\text{refl}}(\mathbf{x})\| = s(\mathbf{x}), \quad \mathbf{x} \in \Omega,$$

$$(8.11) \quad \tau_{\text{refl}}(\mathbf{x}) = \tau_{\text{in}}(\mathbf{x}), \quad \mathbf{x} \in \Gamma,$$

$$(8.12) \quad \left. \frac{\partial\tau_{\text{refl}}}{\partial x} \right|_{\mathbf{x}} = \left. \frac{\partial\tau_{\text{in}}}{\partial x} \right|_{\mathbf{x}}, \quad \mathbf{x} \in \Gamma,$$

$$(8.13) \quad \left. \frac{\partial\tau_{\text{refl}}}{\partial y} \right|_{\mathbf{x}} = - \left. \frac{\partial\tau_{\text{in}}}{\partial y} \right|_{\mathbf{x}}, \quad \mathbf{x} \in \Gamma.$$

Note the minus sign in (8.13). This corresponds to a specular reflection from the “wall”  $\Gamma$ . After we compute  $T_{\text{in}}$  and  $T_{\text{refl}}$  numerically, we can then compute the geometric spreading  $J_{\text{in}}$  and reflected geometric spreading  $J_{\text{refl}}$ . Since the reflecting set is flat, we can use the boundary condition [29]:

$$(8.14) \quad J_{\text{refl}}(\mathbf{x}) = J_{\text{in}}(\mathbf{x}), \quad \mathbf{x} \in \Gamma.$$




FIG. 8.3. Plots comparing domain size ( $|\Omega_h|$ ) and  $\ell_\infty$  and RMS errors for  $T$  and  $\nabla T$ .

Afterwards, we use (7.18) to obtain:

$$(8.15) \quad U(\mathbf{x}) = A_{\text{in}}(\mathbf{x}) \exp(-i\omega T_{\text{in}}(\mathbf{x})) + A_{\text{refl}}(\mathbf{x}) \exp(-i\omega T_{\text{refl}}(\mathbf{x})).$$

In Figure 8.2, we plot  $T_{\text{in}}$ ,  $T_{\text{out}}$ , and the real part of  $U$ .

**8.2. Experimental results.** The results of our numerical experiments evaluating the JMMs described in Section 4 are presented in Table 8.1 and Figures 8.3 and 8.4. The numerical tests for JMM4, which uses cell marching, are given in Ta-




FIG. 8.4. Plots comparing CPU runtime in seconds and errors.

ble 8.2 and Figures 8.5 and 8.6. An example where the geometric spreading and amplitude are computed using cell marching method is shown in Figure 8.7.

For more benign choices of  $s$ , the errors generally convergence with  $O(h^3)$  accuracy for  $T$  and  $O(h^2)$  accuracy for  $\nabla T$  in the RMS error. For the special case of  $s \equiv 1$ , the gradients also converge with nearly  $O(h^3)$  accuracy. For more challenging nonlinear choices of  $s$ , the eikonal converges with somewhere between  $O(h^2)$  and  $O(h^3)$  accuracy, while the gradient converges with nearly  $O(h^2)$  accuracy.

We note that in some cases the gradient begins to diverge for large problem sizes.




FIG. 8.5. Domain size vs. RMS error for JMM4.

	JMM	$E_{\max}(T)$	$E_{\text{RMS}}(T)$	$E_{\max}(\nabla T)$	$E_{\text{RMS}}(\nabla T)$
Constant	#1	2.87	2.87	2.28	2.72
	#2	2.87	2.87	2.28	2.72
	#3	2.87	2.87	2.28	2.72
Linear #1	#1	2.77	2.85	2.14	2.52
	#2	2.77	2.85	1.70	2.48
	#3	2.86	2.87	2.28	2.73
Linear #2	#1	2.48	2.52	1.70	1.97
	#2	2.38	2.52	1.16	1.88
	#3	3.03	3.03	2.70	3.02
Sine	#1	2.76	2.57	1.77	2.09
	#2	2.51	2.46	1.58	1.94
	#3	2.37	2.38	1.54	1.79
Sloth	#1	2.39	2.48	1.49	1.84
	#2	2.37	2.47	0.87	1.73
	#3	2.15	2.21	1.47	1.76

TABLE 8.1

The order of convergence  $p$  for each combination of test problems and solvers, computed for different types of errors and fit as  $Ch^p$ .

This occurs because our tolerance for minimizing  $F$  is not small enough, and also because  $\nabla^2 F$  is  $O(h)$ . For our application, our goal is to save memory and compute time by using a higher-order solver; it is unlikely we would solve problems with such a fine discretization in practice. At the same time, choosing the tolerance for numerical minimization based on  $h$  is of interest—partly to see how much time can be saved for coarser problems, but also to determine to what extent the full order of convergence can be maintained using different floating point precisions.

The JMMs using cubic approximations for  $\varphi$  tend to perform better than those using quadratic approximations when  $s$  is nonlinear and the characteristics of (1.1) are not circular arcs. When  $s$  corresponds to a linear speed of sound, the JMMs with quadratic  $\varphi$  are a suitable choice, generally outperforming the “cubic  $\varphi$ ” solvers, exhibiting cubically (or nearly cubically) convergent RMS errors in  $T$  and  $\nabla T$ . This is a useful finding since the simplified solver requires fewer floating-point operations per update, and since linear speed of sound profiles (e.g., as a function of a linear temperature profile) are a frequently occurring phenomenon in room acoustics.

	$\tau - T$	$\tau_x - T_x$	$\tau_y - T_y$	$\tau_{xx} - T_{xx}$	$\tau_{xy} - T_{xy}$	$\tau_{yy} - T_{yy}$
Constant	3.09	3.11	3.11	2.01	2.05	2.01
Linear #1	2.99	2.43	2.40	1.39	2.01	1.39
Linear #2	2.10	1.76	1.72	0.77	1.25	0.77
Sine	2.91	1.80	1.89	0.73	1.31	0.80
Sloth	2.03	1.76	1.75	0.75	1.33	0.76

TABLE 8.2

The order of convergence  $p$  for *JMM4* for each component of the total 2-jet of  $\tau$ , computed from least squares fits of the RMS error. The fits only incorporate the 4th through the 8th problem sizes to avoid artifacts for small and large problem sizes. See Figure 8.5.




FIG. 8.6. A plot of the pointwise convergence at each point in  $\Omega_h$  for *JMM4*. To obtain these plots, starting with  $N = 129$ , we decimate each larger problem size (up to  $N = 2,049$ ) to a  $129 \times 129$  grid, and do a least squares fit at each point. This gives us an estimate of the order of convergence at each point.

**9. Online Package.** To recreate our results, to experiment with these solvers, and to understand their workings, a package has been made available online on GitHub at <https://github.com/sampotter/jmm/tree/jmm-sisc-figures>. Details explaining how to obtain this package and the collect the results are available at this link.

**10. Conclusion.** We have presented a family of semi-Lagrangian label-setting methods (à la the fast marching method) which are high-order and compact, which we refer to as jet marching methods (JMMs). We examine a variety of approaches to formulating one of these solvers, and in 2D, provide extensive numerical results




FIG. 8.7. Plots related to computing the amplitude and a numerical approximation to the solution to  $(\Delta + \omega^2 s(\mathbf{x})^2)u(\mathbf{x}) = \delta(\mathbf{x})$ , denoted  $U(\mathbf{x})$  for the Linear #1 test problem. Left: the geometric spreading. Middle: the amplitude function. Right: the numerical solution  $U$ .

demonstrating the efficacy of these approaches. We show how a form of “adaptive” cell-marching can be done which is compatible with our stencil compactness requirements, although this scheme no longer displays optimal locality.

Our solvers are motivated by problems involving repeatedly solving the eikonal equation in complicated domains where:

- time and memory savings via the use of high-order solvers,
- high-order local knowledge of characteristic directions,
- and compactness of the solver’s “stencil” (the neighborhood over which the semi-Lagrangian updates require information)

is paramount. In particular, our goal is to parametrize the multipath eikonal in a complicated polyhedral domain in a work-efficient manner. This solver is a necessary ingredient for carrying out this task.

We will be continue to work along the following directions:

- Extension to regular grids in 3D, which should be straightforward and yield considerable savings over existing approaches, and extension to unstructured simplex meshes in 2D and 3D. Especially in 3D, this problem is more complicated, requiring the computation of “causal stencils” [22, 41].
- A rigorous proof of convergence for the solvers developed in this work, including a careful investigation of the conditions resulting in cubic convergence for both the eikonal and its gradient as observed in the case of constant and linear speed functions.

**11. Acknowledgments.** This work was partially supported by NSF Career Grant DMS1554907 and MTECH Grant No. 6205. We thank Prof. Ramani Duraiswami for the illuminating discussions and trenchant observations provided throughout the course of this work. ■

## REFERENCES

- [1] P.-A. ABSIL, R. MAHONY, AND R. SEPULCHRE, *Optimization algorithms on matrix manifolds*, Princeton University Press, 2009.
- [2] J. B. ALLEN AND D. A. BERKLEY, *Image method for efficiently simulating small-room acoustics*, *The Journal of the Acoustical Society of America*, 65 (1979), pp. 943–950.
- [3] G. S. ÁVILA AND J. B. KELLER, *The high-frequency asymptotic field of a point source in an inhomogeneous medium*, *Communications on Pure and Applied mathematics*, 16 (1963), pp. 363–381.
- [4] V. M. BABICH AND N. Y. KIRPICHNIKOVA, *The boundary-layer method in diffraction problems*, vol. 3, Springer, 1979.
- [5] J.-D. BENAMOU, *Big ray tracing: Multivalued travel time field computation using viscosity*

- solutions of the eikonal equation*, Journal of Computational Physics, 128 (1996), pp. 463–474.
- [6] J.-D. BENAMOU, *Multivalued solution and viscosity solutions of the eikonal equation*. 1997.
  - [7] J.-D. BENAMOU, *An introduction to eulerian geometrical optics (1992–2002)*, Journal of scientific computing, 19 (2003), pp. 63–93.
  - [8] J.-D. BENAMOU, S. LUO, AND H. ZHAO, *A compact upwind second order scheme for the eikonal equation*, Journal of Computational Mathematics, (2010), pp. 489–516.
  - [9] F. BORNEMANN AND C. RASCH, *Finite-element discretization of static hamilton-jacobi equations based on a local variational principle*, Computing and Visualization in Science, 9 (2006), pp. 57–69.
  - [10] P. BRUNET, *Increasing the smoothness of bicubic spline surfaces*, Computer Aided Geometric Design, 2 (1985), pp. 157–164.
  - [11] M. K. CAMERON, *Jet marching method*, September 2020, <https://youtu.be/Ze9AeDbADVM>.
  - [12] A. CHACON AND A. VLADIMIRSKY, *Fast two-scale methods for eikonal equations*, SIAM Journal on Scientific Computing, 34 (2012), pp. A547–A578.
  - [13] A. CHANDAK, C. LAUTERBACH, M. TAYLOR, Z. REN, AND D. MANOCHA, *Ad-frustum: Adaptive frustum tracing for interactive sound propagation*, IEEE Transactions on Visualization and Computer Graphics, 14 (2008), pp. 1707–1722.
  - [14] D. L. CHOPP, *Some improvements of the fast marching method*, SIAM Journal on Scientific Computing, 23 (2001), pp. 230–244.
  - [15] G. FARIN, *Curves and surfaces for computer-aided geometric design: a practical guide*, Elsevier, 2014.
  - [16] M. S. FLOATER, *Chordal cubic spline interpolation is fourth-order accurate*, IMA Journal of Numerical Analysis, 26 (2006), pp. 25–33.
  - [17] S. FOMEL, S. LUO, AND H. ZHAO, *Fast sweeping method for the factored eikonal equation*, Journal of Computational Physics, 228 (2009), pp. 6440–6455.
  - [18] J. V. GÓMEZ, D. ALVAREZ, S. GARRIDO, AND L. MORENO, *Fast methods for eikonal equations: an experimental survey*, IEEE Access, (2019).
  - [19] A. GRIEWANK AND A. WALTHER, *Evaluating derivatives: principles and techniques of algorithmic differentiation*, vol. 105, Siam, 2008.
  - [20] H. HAGEN AND G. SCHULZE, *Automatic smoothing with geometric surface patches*, Computer Aided Geometric Design, 4 (1987), pp. 231–235.
  - [21] J. B. KELLER, *Geometrical theory of diffraction*, JOSA, 52 (1962), pp. 116–130.
  - [22] R. KIMMEL AND J. A. SETHIAN, *Computing geodesic paths on manifolds*, Proceedings of the national academy of Sciences, 95 (1998), pp. 8431–8435.
  - [23] R. G. KOUYOUMJIAN AND P. H. PATHAK, *A uniform geometrical theory of diffraction for an edge in a perfectly conducting surface*, Proceedings of the IEEE, 62 (1974), pp. 1448–1461.
  - [24] S. LUO, J. QIAN, AND R. BURRIDGE, *High-order factorization based high-order hybrid fast sweeping methods for point-source eikonal equations*, SIAM Journal on Numerical Analysis, 52 (2014), pp. 23–44.
  - [25] J.-C. NAVE, R. R. ROSALES, AND B. SEIBOLD, *A gradient-augmented level set method with an optimally local, coherent advection scheme*, Journal of Computational Physics, 229 (2010), pp. 3802–3827.
  - [26] R. D. NEIDINGER, *Introduction to automatic differentiation and matlab object-oriented programming*, SIAM review, 52 (2010), pp. 545–563.
  - [27] F. E. NICODEMUS, *Directional reflectance and emissivity of an opaque surface*, Applied optics, 4 (1965), pp. 767–775.
  - [28] M. POPOV, I. PŠENČÍK, AND V. ČERVENÝ, *Computation of ray amplitudes in inhomogeneous media with curved interfaces*, Studia Geophysica et Geodaetica, 22 (1978), pp. 248–258.
  - [29] M. M. POPOV, *Ray theory and Gaussian beam method for geophysicists*, EDUFBA, 2002.
  - [30] S. F. POTTER AND M. K. CAMERON, *Ordered line integral methods for solving the eikonal equation*, Journal of Scientific Computing, 81 (2019), pp. 2010–2050.
  - [31] D. QI AND A. VLADIMIRSKY, *Corner cases, singularities, and dynamic factoring*, Journal of Scientific Computing, 79 (2019), pp. 1456–1476.
  - [32] J. QIAN, L. YUAN, Y. LIU, S. LUO, AND R. BURRIDGE, *Babich’s expansion and high-order eulerian asymptotics for point-source helmholtz equations*, Journal of Scientific Computing, 67 (2016), pp. 883–908.
  - [33] N. RAGHUVANSHI AND J. SNYDER, *Parametric wave field coding for precomputed sound propagation*, ACM Transactions on Graphics (TOG), 33 (2014), p. 38.
  - [34] N. RAGHUVANSHI AND J. SNYDER, *Parametric directional coding for precomputed sound propagation*, ACM Transactions on Graphics (TOG), 37 (2018), p. 108.
  - [35] L. SAVIOJA AND U. P. SVENSSON, *Overview of geometrical room acoustic modeling techniques*,

- The Journal of the Acoustical Society of America, 138 (2015), pp. 708–730.
- [36] C. SCHISSLER, R. MEHRA, AND D. MANOCHA, *High-order diffraction and diffuse reflections for interactive sound propagation in large environments*, ACM Transactions on Graphics (TOG), 33 (2014), p. 39.
  - [37] B. SEIBOLD, J.-C. NAVE, AND R. R. ROSALES, *Jet schemes for advection problems*, arXiv preprint arXiv:1101.5374, (2011).
  - [38] J. A. SETHIAN, *A fast marching level set method for monotonically advancing fronts*, Proceedings of the National Academy of Sciences, 93 (1996), pp. 1591–1595.
  - [39] J. A. SETHIAN, *Fast marching methods*, SIAM review, 41 (1999), pp. 199–235.
  - [40] J. A. SETHIAN, *Level set methods and fast marching methods: evolving interfaces in computational geometry, fluid mechanics, computer vision, and materials science*, vol. 3, Cambridge University Press, 1999.
  - [41] J. A. SETHIAN AND A. VLADIMIRSKY, *Fast methods for the Eikonal and related Hamilton–Jacobi equations on unstructured meshes*, Proceedings of the National Academy of Sciences, 97 (2000), pp. 5699–5703.
  - [42] J. A. SETHIAN AND A. VLADIMIRSKY, *Ordered upwind methods for static Hamilton–Jacobi equations: theory and algorithms*, SIAM Journal on Numerical Analysis, 41 (2003), pp. 325–363.
  - [43] G. E. SHILOV AND R. A. SILVERMAN, *Linear Algebra*, Prentice-Hall, 1971.
  - [44] M. M. SLOTNICK, *Lessons in seismic computing: A memorial to the author*, Society of exploration geophysicists, 1959.
  - [45] J. STOER AND R. BULIRSCH, *Introduction to numerical analysis*, vol. 12, Springer Science & Business Media, 2013.
  - [46] J. N. TSITSIKLIS, *Efficient algorithms for globally optimal trajectories*, IEEE Transactions on Automatic Control, 40 (1995), pp. 1528–1538.
  - [47] R. VERSTEEG, *The marmousi experience: Velocity model determination on a synthetic complex data set*, The Leading Edge, 13 (1994), pp. 927–936.
  - [48] T. XIONG, M. ZHANG, Y.-T. ZHANG, AND C.-W. SHU, *Fast sweeping fifth order WENO scheme for static Hamilton–Jacobi equations with accurate boundary treatment*, Journal of Scientific Computing, 45 (2010), pp. 514–536.
  - [49] S. YANG, S. F. POTTER, AND M. K. CAMERON, *Computing the quasipotential for nongradient SDEs in 3D*, Journal of Computational Physics, 379 (2019), pp. 325–350.
  - [50] Y.-T. ZHANG, H.-K. ZHAO, AND J. QIAN, *High order fast sweeping methods for static Hamilton–Jacobi equations*, Journal of Scientific Computing, 29 (2006), pp. 25–56.
  - [51] H. ZHAO, *A fast sweeping method for eikonal equations*, Mathematics of computation, 74 (2005), pp. 603–627.