

VISIBILITY AND ITS DYNAMICS IN A PDE BASED IMPLICIT FRAMEWORK

YEN-HSI RICHARD TSAI, LI-TIEN CHENG, STANLEY OSHER
PAUL BURCHARD AND GUILLERMO SAPIRO

ABSTRACT. We investigate the problem of determining visible regions given a set of (moving) obstacles and a (moving) vantage point. Our approach to this problem is through an implicit framework, where the obstacles are represented by a level set function. The visibility problem is formally formulated as a boundary value problem (BVP) of a first order partial differential equation. It is based on the continuation of values along the given ray field. We propose an efficient one-pass, multi-level algorithm for the construction the solution on the grid. Furthermore, we study the dynamics of shadow boundaries on the surfaces of the obstacles when the vantage point moves with a given trajectory. In all of these situations, topological changes such as merging and breaking occur in the regions of interest. These are automatically handled by the level set framework proposed here. Finally, we obtain additional useful information through simple operations in the level set framework.

1. INTRODUCTION

In this paper, we consider the visibility problem described as follows: given a collection of hypersurfaces representing the boundaries of objects, called the occluders, in two or three dimensional space, determine the regions of space or on the surfaces visible to a given observer. In real world applications, this problem must be solved efficiently. Generalizations of the visibility problem are just as, if not more, important; this includes the case of a moving rather than static observer and the determination of regions visible for all time or invisible for all time in this situation. We begin with the basic visibility problem for simplicity, and address parts of the dynamic problem later on.

The visibility problem can be reformulated into a problem of determining light and dark regions given a point light source. Under this point of view, a more precise set of assumptions we make in the visibility problem include: a space composed

Date: June 2003.

Research of the first and the third author is supported by ONR N00014-97-1-0027, DARPA/NSF VIP grant NSF DMS 9615854 and ARO DAAG 55-98-1-0323

Research for the last author is supported by ONR,NSF, PECASE, and CAREER.

of a homogeneous medium and objects with nonreflecting and nondiffracting surfaces. Furthermore, we disregard interference, assuming that the distances between objects are large compared to the wavelength of light. Under these conditions, light rays travel in straight lines and are obliterated upon contact with the surface of an object. Thus a point is called visible with respect to a vantage point, the observer, if the line segment between the point and the vantage point does not intersect any of the obstructing objects or their surfaces in space.

Even under these simplifying assumptions, the visibility problem arises as a crucial part of numerous applications in different scientific fields, including rendering, visualization [13], etching [1], the modeling of melting ice [5], surveillance, navigation, and inverse problems, to name a few. In the case of computer graphics and rendering, for example, determination of the visible portions of object surfaces allows for those portions alone to be rendered, thus significantly saving costly computation in unnecessary (invisible) regions. In some modeling problems, such as etching and melting ice that we listed above, visibility is used to find out how certain quantities of interest accumulate, given a radiating source and a (dynamic) surface configuration. There are also variational problems that minimize the corresponding energy functionals over the visible regions of the ambient space, see e.g. [14].

Currently there are numerous algorithms for solving the visibility problem using explicit surface representations. For example, the work of [9] and [11] uses linearity to process triangulated surfaces. A detailed review of related work on the visibility problem, especially concerning explicit surfaces, can be found in [10]. Furthermore, there are a variety of visibility algorithms from computational geometry (see, e.g., [2, 3]). These algorithms often combine special data structures and related algorithms for efficient decomposition and information retrieval of the configuration space. Some are even implemented commercially in hardware to accelerate the solution.

While explicit surfaces, for example triangulated surfaces, are used in a majority of computer graphics and vision applications, implicitly represented surfaces are gaining more attention. This is partly due to the fact that in many applications, the data (i.e. surfaces) is obtained originally and naturally in an implicit form. It is also because of the fact that more and more problems are formulated and solved using the level set method [16]. Hence, it is natural to work directly with the implicit data without converting to a different explicit representation. Currently, visibility algorithms for implicit surfaces mostly consist of sending rays out from the vantage

point to a point of interest (or the reverse) and testing for intersections with the surfaces of the objects using information arising from the implicit formulation.

Another idea in determining whether a point is visible to a given observer is to compare the geodesic and Euclidean distances between the observer and that point. See [22] for an example of this approach. The geodesic distance between two points is the distance in the space in the presence of obstacles, namely the objects. Let \mathbf{x} represent the point of interest and \mathbf{x}_o represent the observer point. The geodesic distance can thus be calculated by solving the Eikonal equation

$$H(\phi)|\nabla u| = 1,$$

with condition $u(\mathbf{x}_o) = 0$, and $u = \infty$ inside of the occluders. Here, $H(x) = \chi_{[0,\infty)}(x)$ is the characteristic function of $[0, \infty)$. Thus the point \mathbf{x} is occluded if and only if

$$u(\mathbf{x}) > |\mathbf{x} - \mathbf{x}_o|.$$

However, this algorithm as it was implemented is $O(N \log N)$, where N is the number of grid points (see, e.g., [27]). Furthermore, numerical implementation of the Heaviside function may cause problems for accuracy.

Our proposed method for ray tracing is different from any of the above. In essence, we send out rays in an implicit manner so as to propagate the causality relation of visibility. This implicit framework for visibility offers many other advantages. For example, the visibility information can be interpreted as the solution of simple first order PDEs and [25] offers a near optimal solution method on the grid. The dynamics of the visibility with respect to moving vantage point or dynamic surfaces can be derived and tracked implicitly within the same framework. Our method retains nearly all the benefits of a level set method, including painless Boolean operations of sets, automatic resolution of surfaces as well as the incorporation of geometric information and the handling of various surface topologies afforded. By embedding the visibility in a Lipschitz continuous function, we can obtain much more information. For example, if one is in the shadow and wants to “be seen” as soon as possible, then one can simply follow the gradient of ψ . Moreover, by the continuous nature of our solution ψ , its application to generating “diffused shadows” in graphics is straightforward. Furthermore, using the same framework and the well developed level set calculus and numerics, one can start solving variational problems on the visibility numerically and efficiently [24]. This will then relate to classical “guarding cameras” or “pursuer-evader” problems in computational geometry and robotics.

In our case of visibility, a real valued function ϕ of two or three dimensions, called the level set function, is introduced. The zero level set of this function represents the surfaces of the occluding objects, and the points where ϕ is negative represent the interior of the objects. Several efficient algorithms have been developed in the literature to obtain this representation. There are also methods describing level set implementation of space partitioning schemes such as octrees. Thus our level set method for visibility will use this function ϕ whenever the objects are considered.

In our approach, the visibility problem is first formally formulated as a boundary value problem for a first order PDE whose solution is constructed by correctly extending the boundary values along the ray fields. The solution of this problem is another level set function, which we called ψ . We then introduce a multi-level algorithm for this boundary value problem for a given fixed vantage point. This algorithm constructs the occlusion boundary, the interface separating visible from invisible. At each resolution level, we solve a radially defined causality relation on a given grid *in one pass*, obtaining not only a conservative estimate of the visible and invisible regions but a locally second order approximation of the occlusion boundary. Thus, our algorithm is independent of both the convexity of the occluders and the grid geometry, and its parallelization is straightforward. In comparison with the method using geodesic distance that is described above, our algorithm in a more primitive form is $O(N)$, a factor of $\log(N)$ faster.

In the second part of the paper, we extend our study to the dynamic visibility problem. In this case, we consider a moving vantage point. Obviously the static visibility problem can be applied at each time to solve this problem, and our algorithm can be used to solve it efficiently enough. However, this static approach does not give us other useful information about the dynamics; for instance, how fast a point in space will become visible or invisible. In many cases, the problem can be solved even faster if the visibility at a previous time is used effectively to produce visibility at future time. Thus we study the dynamics of curves on the occluders that separate light and dark regions on the occluders. The curves in fact can be represented using a level set approach, following the work of [6, 7]. We derive motion laws for all these types of curves and evolve them under the level set framework. Thus, this part of our work complements the book of Cipolla and Giblin [8] which discusses the reconstruction of shape from the perspective (orthogonal) projection of the horizons. To complete our study of visibility dynamics, we derive an emergence-time estimate to predict an occluded object's emergence into view. Obviously, researchers in the computer vision community have also studied

similar topics, which are called visual events. However, their assumptions usually involve objects which are explicitly represented as unions of simple geometrical shapes. We stress here again that we provide a framework to work directly with implicit data that are easy to obtain nowadays without having to convert between data representations.

Through out this paper, we use the following notation:

- The space in which we work will be \mathbb{R}^d , where $d = 2$ or 3 .
- \mathbf{x}_o denotes the position of the vantage point, or observer. We further assume that \mathbf{x}_o never lies in the interior of the objects.
- Ω is a set of connected domains whose closure denotes the objects in question. Furthermore, let $\Gamma = \partial\Omega$.
- ϕ denotes the level set function representing the objects of interest. We may further assume that ϕ is the signed distance function to Γ . This particular level set function can be efficiently computed using fast algorithms such as the fast marching method of [27] or fast sweeping methods [26].
- We define the view direction vector pointing from \mathbf{x}_o to \mathbf{x} by $v(\mathbf{x}_o, \mathbf{x}) = (\mathbf{x} - \mathbf{x}_o)/|\mathbf{x} - \mathbf{x}_o|$. When the context is clear, we will drop the arguments and write simply $v(\mathbf{x})$ or v .
- Let \mathbf{x}_1 and \mathbf{x}_2 denote two points in space. We say $\mathbf{x}_1 \preceq \mathbf{x}_2$ (\mathbf{x}_1 is “before” \mathbf{x}_2) if the conditions $v(\mathbf{x}_o, \mathbf{x}_1) = v(\mathbf{x}_o, \mathbf{x}_2)$ and $|\mathbf{x}_1 - \mathbf{x}_o| \leq |\mathbf{x}_2 - \mathbf{x}_o|$ are satisfied. We also define the strict relation \prec if the condition $|\mathbf{x}_1 - \mathbf{x}_o| \leq |\mathbf{x}_2 - \mathbf{x}_o|$ above is replaced by $|\mathbf{x}_1 - \mathbf{x}_o| < |\mathbf{x}_2 - \mathbf{x}_o|$.
- A point $\mathbf{y} \in \Gamma$ is called a horizon point if and only if $v(\mathbf{x}_o, \mathbf{y}) \cdot \mathbf{n}(\mathbf{y}) = 0$, where $\mathbf{n}(\mathbf{y})$ is the outer normal of Γ at \mathbf{y} . The horizon thus refers to the set of horizon points.
- A point $\mathbf{y} \in \Gamma$ is a terminator point if and only if there is a point \mathbf{y}^* such that: 1) $\mathbf{y}^* \prec \mathbf{y}$ and 2) \mathbf{y}^* is a horizon point. The terminator thus refers to the set of terminator points.
- The visible contour refers to the set of visible points of the horizons and terminators.

2. IMPLICIT RAY TRACING

We now set up the foundation of our approach and derive properties of ray tracing of a single point source in an implicit framework. The motivation is that the visibility along each ray emanating from the vantage point satisfies a causality condition: if a point is occluded, then all other points farther away from the vantage

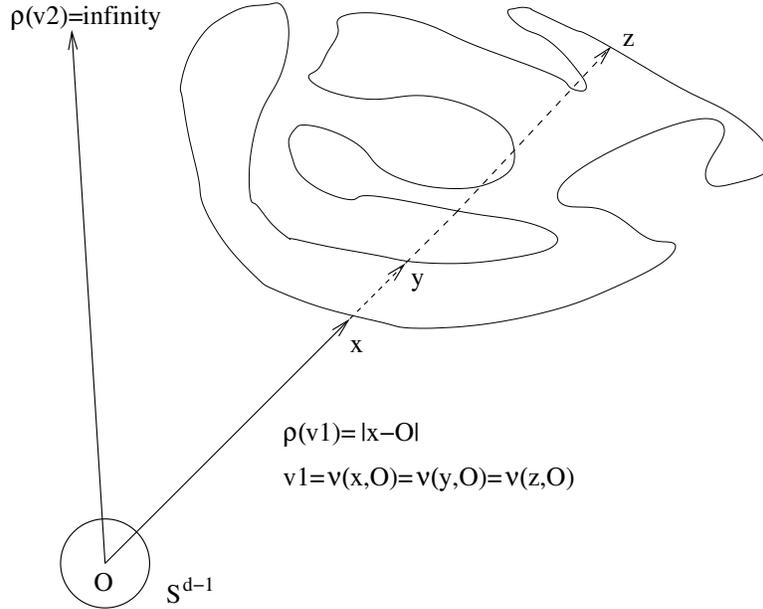


FIGURE 2.1. This figure shows the definition of ρ .

point on the same ray are also occluded, i.e., if \mathbf{x}_1 is occluded and $\mathbf{x}_1 \preceq \mathbf{x}_2$, then \mathbf{x}_2 is also occluded.

We can describe the result of this causality on a sphere centered at the vantage point. Define

$$(2.1) \quad \rho(\mathbf{p}) = \begin{cases} \min_{\mathbf{x} \in \mathbb{R}^d} \{ |\mathbf{x} - \mathbf{x}_0| : v(\mathbf{x}_0, \mathbf{x}) = \mathbf{p}, \phi(\mathbf{x}) \leq 0 \} & \text{if exists} \\ \infty & \text{otherwise} \end{cases}$$

Any given point \mathbf{x} is invisible if $\rho(v(\mathbf{x}, \mathbf{x}_0)) \leq |\mathbf{x} - \mathbf{x}_0|$. Please see figure 2.1 for an example. However, ρ is typically a piecewise continuous function with large jump discontinuity which causes some computational difficulties. In the graphics community, this is closely related to what is called the z-buffer. We defer our discussion of constructing accurate approximations of ρ in a forthcoming paper [23].

In this section, we will provide two closely related interpretations and their corresponding numerical methods. In the first interpretation, the visibility function ψ has a closed analytical expression while in the second interpretation, ψ is the solution of a boundary value problem of a first order linear PDE. The methods to construct approximations to both formulations are closely related and we will extend this type of methods to incorporate a multi-level mesh refinement strategy for efficiency.

2.1. The first interpretation. Our first interpretation of this causality condition is to define a continuous visibility function

$$(2.2) \quad \psi(\mathbf{x}) := \min_{\mathcal{L}(\mathbf{x}_o, \mathbf{x})} \phi(\xi),$$

where $\mathcal{L}(\mathbf{x}_o, \mathbf{x})$ is the line segment connecting \mathbf{x}_o and \mathbf{x} . Thus if $\psi(\mathbf{x})$ is negative, then \mathbf{x} is occluded. We can approximate the value of $\psi(\mathbf{x})$ by $\psi^h(\mathbf{x})$ as follows

$$(2.3) \quad \psi^h(\mathbf{x}) = \min(\psi^h(\mathbf{x}'), \phi(\mathbf{x}))$$

where \mathbf{x}' is some point “immediately before” \mathbf{x} in the ray direction. Here \mathbf{x}' depends on the given grid structure. We will discuss how \mathbf{x}' is defined on a Cartesian grid. As long as the values of $\psi(\mathbf{x}')$ are computed ahead of the computation of $\psi(\mathbf{x})$, this algorithm will be valid. The constraint on the set of updated grid points, \mathbb{X} , can be posed that the unions of the cells whose vertices are all in \mathbb{X} forms star-shaped with respect to the vantage point. For example, consider the case where the vantage point lies on the origin of the coordinate system. The grid points in the first quadrant can be updated row by row, starting from the positive part of the x -axis, in an increasing order of their y -coordinate component. The grid points in each row are updated from left to right, in increasing order of their x -coordinate. We can thus generalize this approach: in 3D, we consider the vantage point as the origin and approximate ψ in each octant separately, and inside each quadrant, we employ similar approach to what is described above.

We write down our basic algorithm as follows:

Algorithm 2.1. (*Basic visibility sweeping*)

1. Set $\psi(\mathbf{x}_o) = \phi(\mathbf{x}_o)$.
2. Do a star-shaped¹ updating sequence on the grid.
3. For each grid point \mathbf{x} , choose \mathbf{x}' depending on the grid geometry.
4. Compute the value of $\psi(\mathbf{x})$ via (2.3).

For clarity of exposition, we defer a detailed description to the next subsection, and also the Appendix. Figure 2.2 shows what ψ should look like in a one space dimension setting.

Before we move on to our multi-level implementation of the above algorithm, we discuss some additional properties of our proposed visibility representation, ψ , and its numerical approximation ψ^h .

Lemma 2.2. *If $|\phi(\mathbf{x}) - \phi(\mathbf{y})| \leq L|\mathbf{x} - \mathbf{y}|$, then $|\psi(\mathbf{x}) - \psi(\mathbf{y})| \leq L|\mathbf{x} - \mathbf{y}|$.*

¹See 5.2

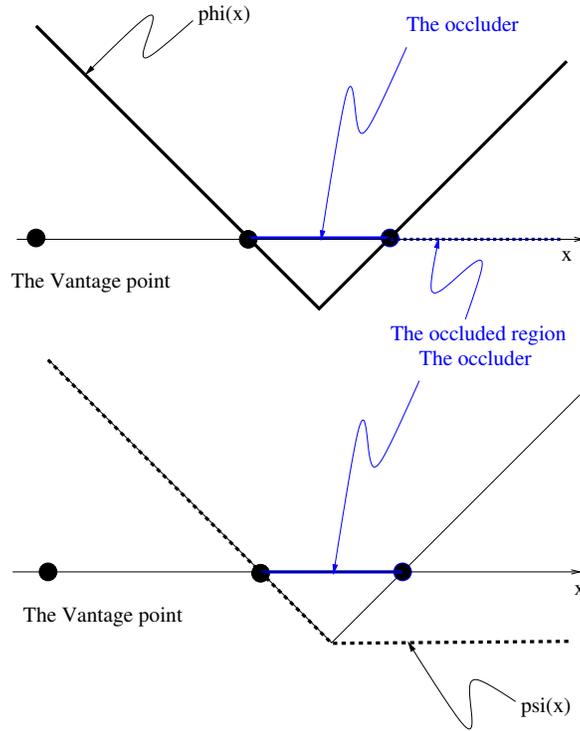


FIGURE 2.2. A demonstration of the motivation of our implicit ray tracing algorithm in one dimension.

Proof. This can be shown directly from the definition:

$$\begin{aligned}
 |\psi(\mathbf{x}) - \psi(\mathbf{y})| &= \left| \min_{\mathcal{L}(\mathbf{x}, \mathbf{x}_o)} \phi - \min_{\mathcal{L}(\mathbf{y}, \mathbf{x}_o)} \phi \right| \\
 &\leq |(\phi(\mathbf{x}_o) - L|\mathbf{x} - \mathbf{x}_o|) - (\phi(\mathbf{x}_o) + L|\mathbf{y} - \mathbf{x}_o|)| \\
 &\leq L|\mathbf{x} - \mathbf{y}|,
 \end{aligned}$$

□

2.2. PDE interpretation and extensions. Here, we propose an approximation scheme to (2.2) using a linear PDE and a thresholding procedure.

Let $r(\mathbf{x})$ be a suitable vector field in the domain Ω such that all the flow lines of r in Ω emanate from the set $O \subsetneq \Omega$. For example, our main focus in this paper will be the case in which $r(\mathbf{x}) = v(\mathbf{x}) = (\mathbf{x} - \mathbf{x}_o)/|\mathbf{x} - \mathbf{x}_o|$, and $O = \{\mathbf{x}_o\}$. In this slightly more general setting, formula (2.2) is generalized to

$$\psi(\mathbf{x}) = \min_{\mathbf{y} \in \tilde{\mathcal{L}}(O, \mathbf{x})} \phi(\mathbf{y}),$$

where $\tilde{\mathcal{L}}(O, \mathbf{x}) = \{\text{the flow line connecting } \mathbf{x}_o \text{ to } O\}$. Defining $\tilde{h}(\mathbf{x}) = \nabla\phi(\mathbf{x}) \cdot r(\mathbf{x})$, we observe that its zeros correspond to the critical point of ϕ along the flow lines that pass through; i.e. $\psi(\mathbf{x}) = \phi(\tilde{\mathbf{y}})$ for some $\tilde{\mathbf{y}} \in \tilde{\mathcal{L}}(O, \mathbf{x})$ and $\nabla\phi(\tilde{\mathbf{y}}) \cdot r(\tilde{\mathbf{y}}) = 0$. Even though our implicit algorithm will not make use of this feature, we will see that \tilde{h} will be used in a later section for identifying a type of important curve on the occluders that we call the horizons.

Our algorithm is the following: Set $\psi|_O = \phi|_O$. At grid point $\mathbf{x}_{i,j}$, $\psi_{i,j}$ is computed by:

- (1) Solve for $\psi_{i,j}$ by upwinding:

$$(2.4) \quad \nabla\psi_{i,j} \cdot r(\mathbf{x}_{i,j}) = 0.$$

- (2) Update

$$\psi_{i,j} = \min(\phi_{i,j}, \psi_{i,j}).$$

In general, the above equation (2.9) can be solved by the fast sweeping method [26] on Cartesian grid. Alternatively, due to the special structure of the characteristics, we proposed an efficient one-pass sweeping algorithm. For simplicity, we will assume that $\mathbf{x}_o = (0, 0)$, $\mathbf{x} = (x, y)$, $r(\mathbf{x}) = (r_1, r_2)$, and shall introduce an upwind scheme in the first quadrant $\Omega_{+,+} = \{x \geq 0, y \geq 0, xy \neq 0\}$. In this region, equation (2.4) is discretized by regular upwind scheme

$$r_1(\mathbf{x}_{i,j})D_x^- \psi_{i,j} + r_2(\mathbf{x}_{i,j})D_y^- \psi_{i,j} = 0,$$

leading to an update formula for $\psi_{i,j}$:

$$(2.5) \quad \psi_{i,j}^h = \left(\frac{1}{r_1(\mathbf{x}_{i,j}) + r_2(\mathbf{x}_{i,j})} \right)^{-1} \left(r_1(\mathbf{x}_{i,j})\psi_{i-1,j}^h + r_2(\mathbf{x}_{i,j})\psi_{i,j-1}^h \right)$$

$$(2.6) \quad := \min(G_{+,+}(\psi_{i-1,j}^h, \psi_{i,j-1}^h), \phi_{i,j}).$$

With this scheme, we compute $\psi_{i,j}^h$ in $\Omega_{+,+}$ in the following order:

Algorithm 2.3. (*Updating scheme for region $\Omega_{+,+}$*)

for $i=0:n$

for $j=0:m$ ($ij \neq 0$)

$$\psi_{i,j}^h = \min(G_{+,+}(\psi_{i-1,j}^h, \psi_{i,j-1}^h), \phi_{i,j}).$$

In fact, the above algorithm is an alternative way of solving (2.2). For $\mathbf{x}_o \in [0, \Delta x] \times [0, \Delta y]$ not lying on a grid point, we use a slightly different discretization. For example for $\mathbf{x}_{1,j}$, $1 \geq j \geq n$,

$$(2.7) \quad D_y^- \psi_{i,j} = -\frac{r_1(\mathbf{x}_{i,j})}{r_2(\mathbf{x}_{i,j})} D_x^- \psi_{i,j-1}.$$

This corresponds to the interpolations described in the Appendix for (2.3). Figure 2.3² provide a result of the above algorithm applied to a real city model. In the general cases, where the integral curves of $r(\mathbf{x})$ are more complicated than straight lines, similar convergence analysis applies, since characteristics do not cross. See Figure 2.4 for a result under this kind of ray field. We may use higher order approximation schemes to obtain more accurate solutions. In Ω^x we can, for instance, use linear multistep methods for advance in the x -direction, and higher order ENO type discretization, see [12], for the computation of ψ_y .

An alternative interpretation of the causality condition stated in the beginning of this section is to formulate it as a boundary value problem for the following first order PDE:

$$(2.8) \quad \nabla\psi \cdot \frac{\mathbf{x} - \mathbf{x}_o}{|\mathbf{x} - \mathbf{x}_o|} = \min(H(\psi - \phi)\nabla\phi \cdot \frac{\mathbf{x} - \mathbf{x}_o}{|\mathbf{x} - \mathbf{x}_o|}, 0), \quad \psi(\mathbf{x}_o) = \phi(\mathbf{x}_o),$$

where $H(z) = \chi_{[0, \infty)}(z)$ is the characteristic function of $[0, \infty)$. More generally, we can handle the situation in which the light rays are not straight lines. Assuming $\psi > 0$ on a set $\Gamma \subset \omega$ and $r(\mathbf{x})$ be a differentiable vector fields such that for the integral curves starting from every interior point of Ω traces back to Γ . The ‘‘generalized’’ visibility information on this vector field can be obtained from ψ that solves

$$(2.9) \quad \nabla\psi \cdot r(\mathbf{x}) = \min(H(\psi - \phi)\nabla\phi \cdot r(\mathbf{x}), 0), \quad \psi(\mathbf{y}) = \phi(\mathbf{y}) \text{ for every } \mathbf{y} \in \Gamma.$$

Consequently, we have an the a priori estimate

$$|\nabla\psi| = \max_{|r|=1} \nabla\psi \cdot r \leq \max_{|r|=1} \nabla\phi \cdot r = |\nabla\phi| \leq L;$$

i.e. Lemma 2.2 holds for solutions of (2.9). We point out, however, that much study is needed for this kind of PDEs with discontinuous coefficients.

2.3. Estimates of visibility change using Lipschitz constant. Our multi-level approach to the visibility problem requires the skipping of large regions which are determined a priori to be either visible or invisible. This hinges upon the ability of determining whether any given voxel is completely ‘‘inside’’ or ‘‘outside’’ of the objects. This can be done conservatively with the help of the Lipschitz constant of the embedding level set function ϕ .

If we assume that we know the values $\phi(\mathbf{y}), \phi(\mathbf{z}) > 0$, and let \mathbf{x} be a point on the line segment joining \mathbf{y} and \mathbf{z} . We would like to know if $\phi(\mathbf{x})$ can be negative. Let C be the Lipschitz constant of ϕ . It dictates the maximal decrease in values from \mathbf{y}

²The authors thank Prof. John Steinhoff and Yonghu ‘‘Tiger’’ Wenren for their help in obtaining this data.

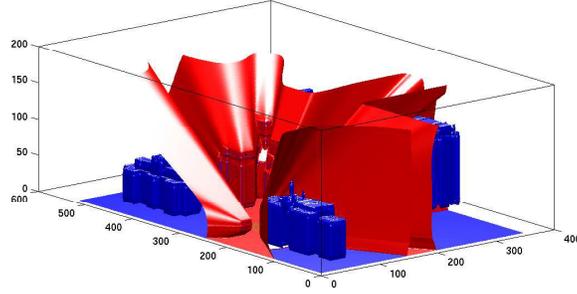


FIGURE 2.3. A visibility result applied to a city grid. The blue surface represents the occluders, and the red surface represents the boundary between visible and invisible regions with respect to the vantage point at the position indicated by the green diamond.

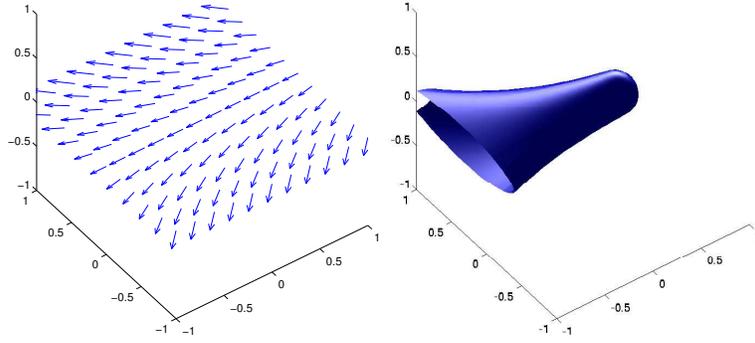


FIGURE 2.4. Visibility under a bending ray field.

to \mathbf{x} to be $L|\mathbf{y} - \mathbf{z}|$, therefore,

$$\phi(\mathbf{x}) \geq \phi(\mathbf{y}) - L|\mathbf{y} - \mathbf{x}|$$

and, similarly

$$\phi(\mathbf{x}) \geq \phi(\mathbf{z}) - L|\mathbf{z} - \mathbf{x}|.$$

Hence, if $0 \leq \phi(\mathbf{y}) - L|\mathbf{y} - \mathbf{x}|$ and $0 \leq \phi(\mathbf{z}) - L|\mathbf{z} - \mathbf{x}|$, we can conclude that the values of ϕ along the line segment connecting \mathbf{y} and \mathbf{z} always stay above zero.

We thus generalize the above observation to determine the sign of ϕ or ψ over a rectangular or cubic region that we call a voxel. Let \mathbf{x}_c be the center point and \mathbf{x}_i the vertices of the given voxel V . If

$$(2.10) \quad \phi(\mathbf{x}_i) + L|\mathbf{x}_c - \mathbf{x}_i| < 0 \quad \forall i,$$

then we know $\phi|_V < 0$ ($V \subset \{\phi < 0\}$). Conversely, if

$$(2.11) \quad \phi(\mathbf{x}_i) + L|\mathbf{x}_c - \mathbf{x}_i| > 0 \quad \forall i,$$

then we know $\phi|_V > 0$. Since ϕ is the signed distance function, as we pointed out in the previous subsection, a Lipschitz constant of ψ can be taken to be 1.

If we can obtain an L_∞ error bound $\mathcal{E}(V; \psi^h)$ of ϕ^h on the give voxel V , we can conservatively estimate both the visible and invisible regions from our approximation ψ^h . I.e.

Case 1. If $(\psi^h(\mathbf{x}_i) + \mathcal{E}(V; \psi^h)) + L|\mathbf{x}_c - \mathbf{x}_i| < 0 \quad \forall i$, then $\psi|_V < 0$. (V is definitely invisible.)

Case 2. If $(\psi^h(\mathbf{x}_i) - \mathcal{E}(V; \psi^h)) + L|\mathbf{x}_c - \mathbf{x}_i| > 0 \quad \forall i$, then $\psi|_V > 0$. (V is definitely visible.)

2.4. A Multi-resolution algorithm. We offer a mesh refinement strategy to further accelerate our one-pass implicit ray tracing algorithms. Essentially, these are all applications of methods of characteristics for first order PDEs.

Assuming the ability of determining whether a given voxel lies completely inside or outside of an occluder³, we propagate these voxels along the rays from the vantage point, and obtained a set of voxels over which the visibility status does not change. We then refine our grid over the remaining region. This is in similar spirit to what is reported in [21].

Given a grid at resolution $2h$, we use \mathcal{V}_{2h}^+ and \mathcal{V}_{2h}^- to represent the subsets of Ω over which the analytic solution ψ is determined to remain positive and negative respectively. Furthermore, let $\mathcal{V}_{2h} = \mathcal{V}_{2h}^+ \cup \mathcal{V}_{2h}^-$ and let $\partial\mathcal{V}_{2h}^\pm$ denote that boundary of \mathcal{V}_{2h}^\pm . We refine the grid over the region $\Omega_h = \Omega_{2h} \setminus \mathcal{V}_{2h}$ and subsequently determine the sets $\mathcal{V}_{h/2}^\pm$.

One way to find $\mathcal{V}_{h/2}^\pm$ is to solve ψ^h for the same PDE with different boundary conditions:

$$(2.12) \quad \begin{cases} \nabla\psi^h \cdot r(x) = \min(H(\psi - \phi)\nabla\phi \cdot r(x), 0) \text{ for } x \in \Omega_h^{(0)}, \text{ the interior of } \Omega_h, \\ BC : \psi^h(x) = \begin{cases} \phi(x) & \text{if } x \in \partial\mathcal{V}_{2h}^+ \text{ and } r(x) \cdot n_\Omega(x) > 0 \\ \psi_{\text{int}}^{2h}(x) & \text{if } x \in \partial\mathcal{V}_{2h}^- \text{ and } r(x) \cdot n_\Omega(x) > 0 \end{cases} \end{cases}$$

where n_Ω is the inner normal of $\partial\Omega_h$, and ψ_{int}^{2h} is a linear interpolant of the grid function ψ^{2h} . We then determine \mathcal{V}_h using the analysis shown in the previous section.

³Later on, we will call these voxels the ‘‘source’’ voxels.

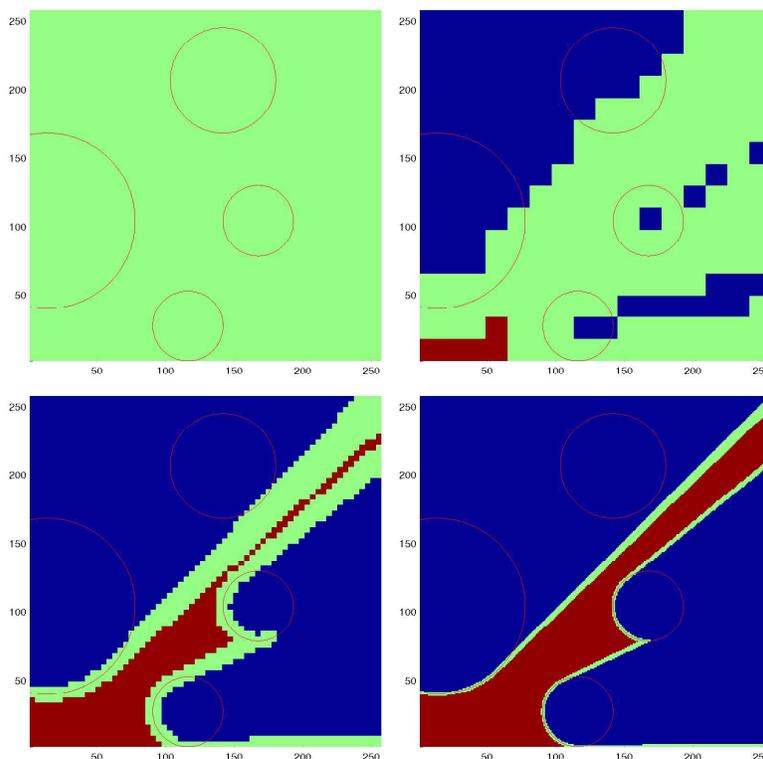


FIGURE 2.5. This is a schematic diagram for the multi-resolution algorithm. Occluded voxels are depicted in blue and visible ones in red. The regions are target for next level refinement, to be refined. The red curves represent the boundaries of the occluders, and the vantage point is positioned at $(1, 1)$. The sizes of the voxels are: 64×64 , 16×16 , 4×4 , and 1×1 .

Hence we can repeat this procedure until the desired resolution is reached. This approach, however, requires either *a priori* or *a posteriori* estimates of the error to the viscosity solution of the given problem. It will be reported in a forthcoming paper by the authors.

The second way is the following: At each resolution level, we construct another function M such that for each voxel V in the domain, $M_h(x) = (\tilde{V}, \sigma, \bar{\phi})$, where the first component \tilde{V} is the “source” voxel of V (see Footnote 3), the second component is the visibility of V ($\sigma = 1$ if visible, -1 if invisible, 0 inconclusive), and the third will be the constant continuation of the values of ϕ along $r(x)$ from the “source” voxels.

We first identify those V_j on which the embedding level set function ϕ is negative; these are the “source” voxels. Set $M(x) = (V_j, -1, \bar{\phi})$. Find V_o that contains

\mathbf{x}_o . If ϕ is positive over V_o , set $M(x) = (V_o, 1, \bar{\phi})$. Let $\Omega_h = \Omega_{2h} \setminus ((\cup_j V_j) \cup V_o)$. We then “solve” the following problem by method of characteristics:

$$(2.13) \quad \begin{cases} \vec{\nabla} M_h \cdot r(x) = 0 & \text{for } x \in \Omega_h^{(0)} \\ BC : M_h(x) = M_{2h}(x) & \text{if } x \in \partial \mathcal{V}_{2h}^+ \cup \partial \mathcal{V}_{2h}^- \text{ and } r(x) \cdot n_\Omega(x) > 0 \end{cases}$$

In a 2D setting, at each vertex x of V , $M_h(x) = \theta M(x_j) + (1 - \theta)M(x_k)$ for some upwind neighbors x_j and x_k , and some $\theta \in [0, 1]$ determined by $r(x)$. However, we modified $M_h(x)$ after the update formula by rounding the second component of M_h that is neither 1 nor -1 to 0. Finally, a voxel V is determined to be in \mathcal{V}_h^- , or \mathcal{V}_h^+ , if, on V , the second component of M_h is -1 , or 1, and the voxels referred to by M_h are all immediate neighbors to each other. Please see Figure 2.5 for a demonstration of this algorithm.

As for complexity, the operation count for our multi-level algorithms is $O(N^{d-1} \log N)$. Here $N = 1/h$ where h is the smallest spatial stepsize used in the multiresolution framework and d is the dimension of the space. The N^{d-1} part of the complexity comes from the fact that a codimension one hypersurface in d dimensional space is being generated under fast sweeping and the $\log N$ part comes from multiresolution. The memory allocation of our algorithm is also $O(N^{d-1} \log N)$, with the $\log N$ part once again due to multiresolution.

We point out here that for applications in which the occluders are triangulated and one is only interested in the visibility information projected on an image plane, there are already many specialized algorithms and hardware designed available. The purpose of our algorithm is to work with implicit data and find visibility information in the whole ambient space without the costly operations of changing representations.

2.5. Multi-scale considerations. If we consider the visibility problem in applications related to human vision, such as 3D virtual environment rendering, it is natural to put a scale parameter into the size of the objects related to the distance of the object from the vantage point. We want to ignore certain **isolated and small** objects that are **far away** from the vantage point using this information. It is important to notice that a collection of closely positioned small objects can form a visible ensemble, seen for example in clouds and trees.

Using the level set representation of the virtual environment in conjunction with the solution properties of certain PDEs, we are able to deal with this issue easily without explicitly considering each object separately. The idea is to *dilate* the interface first so that small objects can merge to form ensembles of larger size. We then

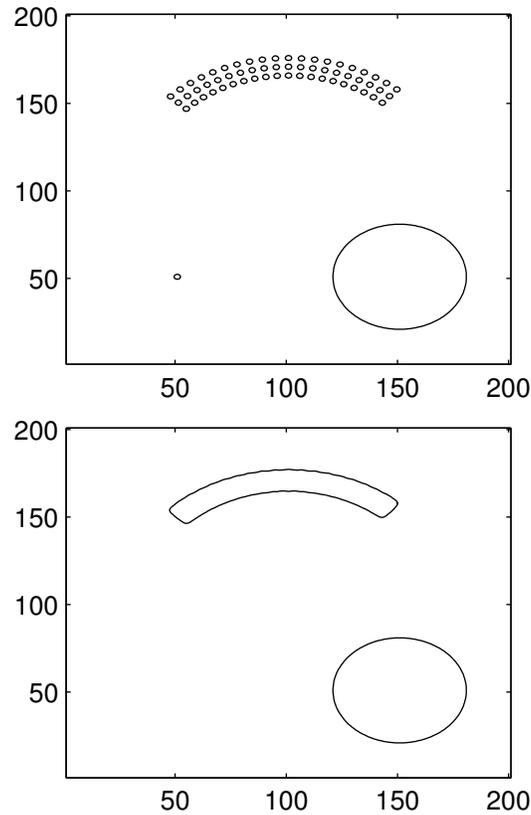


FIGURE 2.6. An example of grouping

shrink the interfaces (one possibility is to perform curvature driven motion) such that remaining small objects will disappear. The result of this approach follows the regularization effect of viscosity solution theory for Hamilton-Jacobi Equations. It is basic mathematical morphology, and can be done easily, see e.g. [4][20].

We shall return to a brief discussion of scales at the end of this paper, after we discuss the dynamics of the visibility.

3. DYNAMIC VISIBILITY

In this section, we introduce an implicit framework for tracking the change in visibility when the vantage point or the underlying occluders are changing in time. Typically, disconnected components of the invisible regions may merge and one single connected component may break up into several pieces; completely occluded objects may emerge to the scene and occlude parts of the domain that

was previously visible. In many of these situations, the topology of the occlusion changes and implicit methods such as the level set methods become attractive options.

For simplicity, we consider the case in which the vantage point is moving and the occluders are stationary. We formulate the visibility problem so that the points which are on the boundaries of the visible regions on the surfaces of the occluders can easily be identified. The dynamics of these points are derived so that one can track the visible regions according to the motion of the vantage point \mathbf{x}_0 .

For a single convex object, the horizon determines the visibility information on the surface. Therefore, tracking the motion of the horizon for all time gives us incremental information on the change of the visible portion of the object. For non-convex objects or multiple objects, these horizon extends the rays to other parts of the surfaces, and thereby creating another type of occlusion boundaries which we coin “terminators”, based on The Merriam-Webster Dictionary. Similar to the horizons, one side of a terminator is visible while the other not. In summary, the points forming the boundaries of visible regions on given surfaces can be placed into two categories:

- points that are part of the horizon;
- points that border *shadows cast by some surface* (terminator).

We shall see that the motion of the horizon is characterized by the orthogonality constraint and it, in turn, becomes a part of the constraints of the terminator motion.

In our level set formulation, we create a continuous function whose zero level set captures the curves described above. We also propose a method that relates each point on the terminator to a point on the horizon of the surface casting the shadow. This description should be global, that is, quantities should vary “continuously” with respect to points not on the surface.

For a single convex object, tracking the horizon is certainly the optimal solution. With the fast sweeping algorithms and local storage strategies, the complexity of the level set approach to track the horizon and terminator curves is formally $O(N)$ both in operation count and in storage. Here, the number N is the number of points used to resolve the curves. In more realistic applications, we need to consider the situations in which 1) the velocity fields for the horizons and terminators become singular, 2) some “hidden” object may appear and create new terminators. Indeed, these are among the most difficult problems in this topic. We will address this difficulty in a later subsection. If the visibility information in the whole ambient space is needed, we think that our previous “static” approach might generate more

elegant and easier-to-implement algorithms with comparable performance. However, the additional information about how shadow boundary moves may be used together with the static visibility algorithm for many dynamic type applications.

In the following, we will first proposed a fast implicit characterization of the horizons and their terminators. The procedures involved consist of our “static” algorithm and some simple boolean operations on sets.

3.1. Finding the horizon and the terminator implicitly.

Finding the horizon. We extend the orthogonality condition that defines the horizon and arrive at

$$(3.1) \quad \bar{h}(\mathbf{x}, t) = (\mathbf{x} - \mathbf{x}_0) \cdot \nabla \phi(\mathbf{x}).$$

Earlier, we have seen that the zeros of this function correspond to the critical point of ϕ along the rays that pass through. We also noted above that \bar{h} completely determines the visibility of any convex object embedded in ϕ :

$$\{\bar{h}(\mathbf{x}) \leq 0\} \cap \{\phi = 0\} \iff \text{visible.}$$

In general cases, where there are multiple objects (convex and nonconvex), \bar{h} does not give exact visibility information anymore. It just provides local visibility information just as local extrema may not be absolute extrema.

Due to its definition (3.1), \bar{h} will still be non-negative on the parts of the surface facing the source, even though those parts are completely occluded. Instead, \bar{h} gives a conservative “estimate” of the shadow:

$$\{\bar{h}(\mathbf{x}) > 0\} \cap \{\phi = 0\} \implies \text{invisible.}$$

Thus the *visible* horizon is

$$\{\phi = 0\} \cap \{\bar{h} = 0\} \cap \{\psi \geq 0\},$$

where ψ is the visibility function coming from our static algorithm. Figure 3.1 gives an example of horizons found this way.

Finding the terminator. How do we find the terminator? Our idea is to overshoot each ray that tangents the visible horizons when it hits another part of Γ ; thus the intersections of these rays and Γ correspond to exactly the terminators. This “overshooting” strategy is, of course, implemented by an auxiliary level set function $\tilde{\psi}$, and $\{\tilde{\psi} = 0\}$ will cut through Γ on the terminator, therefore providing an implicit representation of it. Consider $\tilde{\phi} = \max(\phi, -\bar{h})$, then $\{\tilde{\phi} \leq 0\}$ corresponds to the set $\{\bar{h} \geq 0\} \cap \{\phi \leq 0\}$, a set created by “carving off” a neighborhood of the

visible portions the occluders. We notice that the occlusion generated by the set $\{\tilde{h} \geq 0\} \cap \{\phi \leq 0\}$ is the same⁴ as $\{\phi \leq 0\}$. Therefore, we can construct $\tilde{\psi}$ as the solution of (2.8) with $\tilde{\phi}$ on the right hand side, instead of ϕ . Consequently, the shadow boundaries, $\{\tilde{\psi} = 0\}$, extending from the horizons cuts through portions of the original occluders $\{\phi \leq 0\}$, and the intersections corresponds to the terminators. Figure 3.1 shows the operations described above in a simple two circle setting. Additionally, we can even make $\{\tilde{\psi} = 0\}$ perpendicular to Γ locally around the intersection by iterating on the following PDE used in [15]:

$$\tilde{\psi}_\tau + \text{sgn}(\phi) \nabla \tilde{\psi} \cdot \frac{\nabla \phi}{|\nabla \phi|} = 0.$$

With these characterizations, we can easily identify the visible contours. See Figures 3.2, 3.3, and Figures 3.4, 3.5⁵ for examples. In these figures, the visible portions of the horizons and the terminators are depicted as cyan and yellow curves respectively. A green circle is drawn to reveal the location of the vantage point in each setting. The boundaries between visible and invisible regions are represented by blue surfaces. We observe that the blue surfaces cut through the objects exactly at the visible contours.

3.2. The dynamics of the horizon. Let $\mathbf{x}_o(t)$ be the position of the vantage point and $\mathbf{x}(t)$ be a corresponding point on the horizon at time t . We first consider a single convex occluder Ω embedded by the signed distance function ϕ . Let $n(x)$ denote the outer normal of $\partial\Omega$ at \mathbf{x} . This translates into the following constraints on $\mathbf{x}(t)$:

$$(3.2) \quad \begin{cases} \phi(\mathbf{x}(t)) = 0 \\ (\mathbf{x} - \mathbf{x}_o) \cdot \nabla \phi(\mathbf{x}) = 0 \end{cases}$$

In two dimensions, we can invert the above constraints and derive that

$$(3.3) \quad \dot{\mathbf{x}} = \begin{pmatrix} \dot{x} \\ \dot{y} \end{pmatrix} = \frac{1}{\kappa} \frac{\dot{\mathbf{x}}_o \cdot n(\mathbf{x})}{|\mathbf{x} - \mathbf{x}_o|} \mathbf{v}(\mathbf{x}).$$

Here, κ is the curvature of the occluding surface at \mathbf{x} . In three space dimensions, the horizon becomes a closed curve $\Gamma(s) = \mathbf{x}(s, t)$, where s is the arc length of $\Gamma(s)$. Let P be the plane tangent to \dot{x}_o , passing through $\Gamma(s)$ and x_o . Let $\beta(\sigma)$ be the curve on the intersection of P and $\partial\Omega$. Then, locally at t and x , we have a two dimensional visibility problem on the plane P , in which $\beta(\sigma)$ defines the boundary

⁴modulo a small subset of $\{\phi \leq 0\}$, which we know is invisible by definition.

⁵The terrain data is obtained from <ftp://ftp.research.microsoft.com/users/hhoppe/data/gcanyon/>.

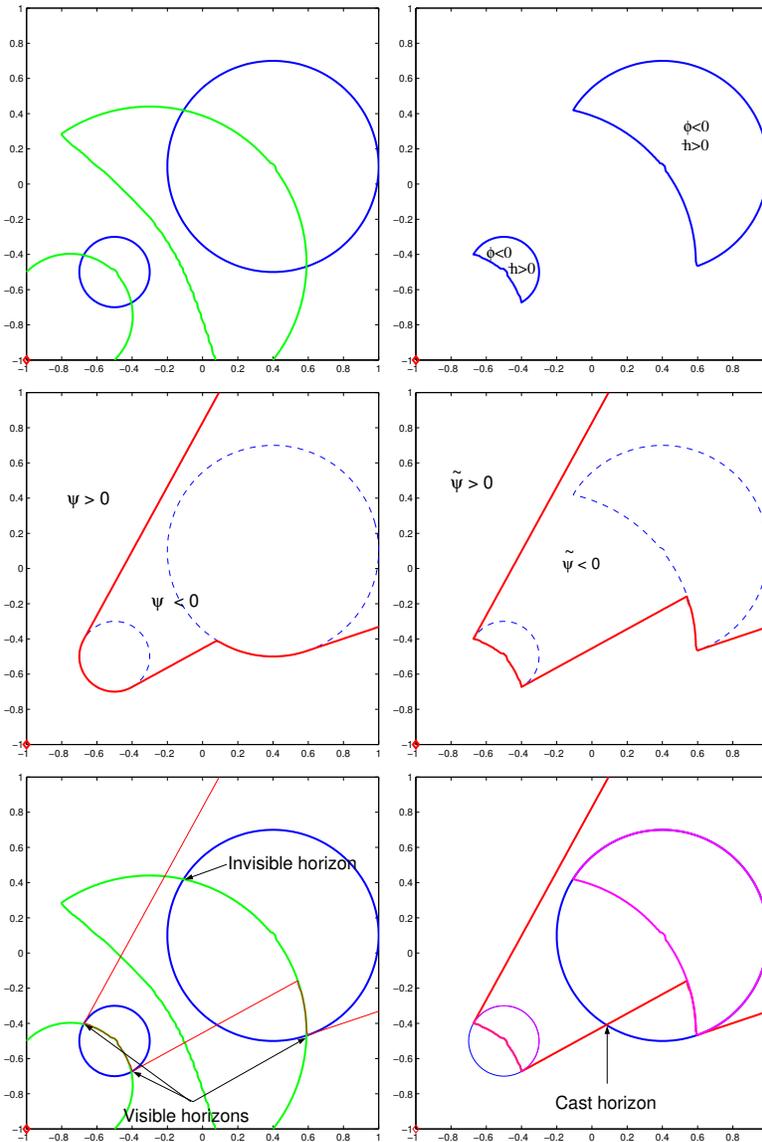


FIGURE 3.1. Finding the visible silhouettes and their casts. The occluders are the two circles depicted by the blue curves, and the vantage point is located at $(-1, -1)$. The green curves are the zero level set of \tilde{h} . Visible silhouettes and their casts are characterized by the intersections of different level set functions as described in the text.

of the objects. Following this reasoning, κ should naturally be taken from $\beta(\sigma)$. See Figure 3.7.

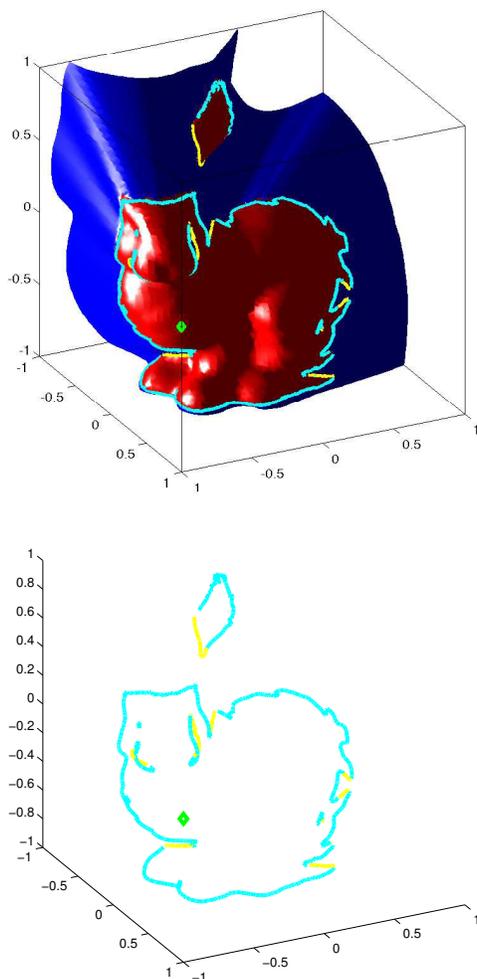


FIGURE 3.2. Visible contour (portions of silhouette and terminator that are visible)

Alternatively and more naturally under our level set formulation, we rederive the above motion law as

$$(3.4) \quad \dot{x} = \frac{II(x-x_0)}{|II(x-x_0)|^2} \left(\dot{x}_0 \cdot \frac{\nabla\phi}{|\nabla\phi|} \right),$$

where II is the *second fundamental form*, which can conveniently be extended to the other level sets and takes the form:

$$II = \frac{1}{|\nabla\phi|} P_{\nabla\phi} \nabla^2 \phi P_{\nabla\phi}.$$

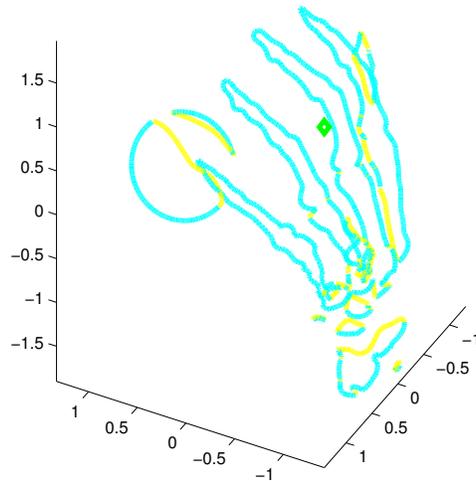
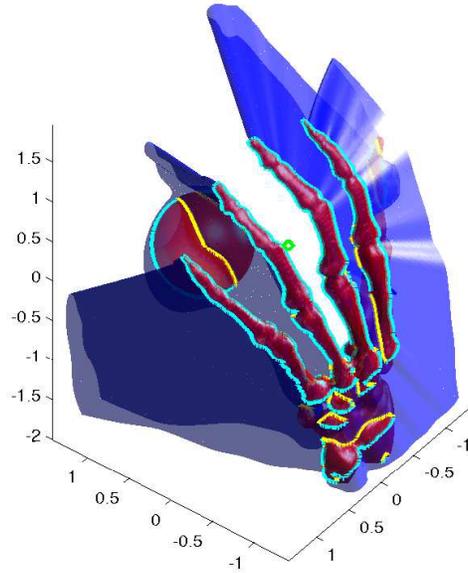


FIGURE 3.3. Visible contour (portions of silhouette and terminator that are visible)

Here, $P_{\nabla\phi}$ is the orthogonal projection matrix projecting vectors to the plane with normal vector parallel to $\nabla\phi$. Thus the κ in (3.3) denotes the normal curvature of the surface in the viewing direction.

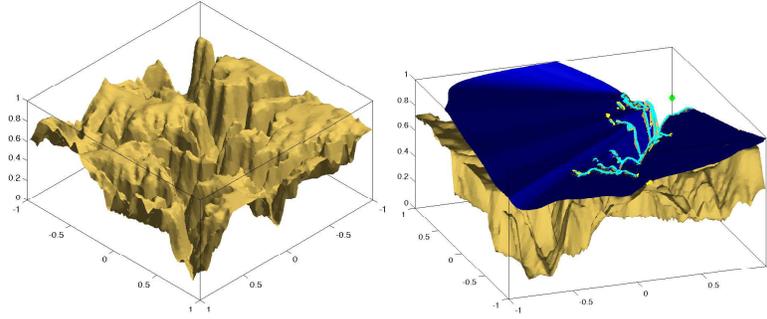


FIGURE 3.4. silhouettes and terminators obtained from the elevation data of Grand Canyon.

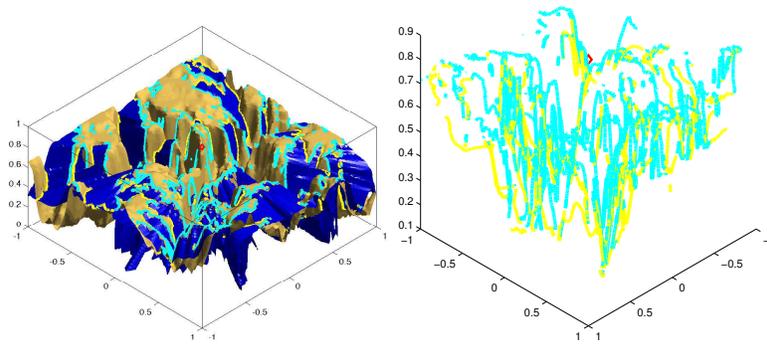


FIGURE 3.5. terminators and terminators obtained from the elevation data of Grand Canyon.

For a detailed derivation and implementation, please see sections 5.4, 5.3 and 5.6. Figure 3.8 shows a result of horizon motion on a nonconvex body.

3.3. The dynamics of the terminator. Assume that \mathbf{x} is a terminator point and $\mathbf{x}^*(\mathbf{x})$ is its generator. In two dimensions, the motion of \mathbf{x} is determined by the following constraints:

$$(3.5) \quad \begin{cases} \phi(\mathbf{x}) = 0, \\ \frac{\mathbf{x} - \mathbf{x}_o}{|\mathbf{x} - \mathbf{x}_o|} = \frac{\mathbf{x}^* - \mathbf{x}_o}{|\mathbf{x}^* - \mathbf{x}_o|}. \end{cases}$$

Inverting, we find that the motion of the terminator can be written as follows:

$$(3.6) \quad \dot{\mathbf{x}} = \frac{1}{\mathbf{v} \cdot \mathbf{n}(\mathbf{x})} \left(\frac{|\mathbf{r}|}{|\mathbf{r}^*|} \mathbf{r}^* \cdot (\mathbf{v}^*)^\perp + \dot{\mathbf{x}}_o \cdot \mathbf{v}^\perp \right) \mathbf{n}^\perp(\mathbf{x}),$$

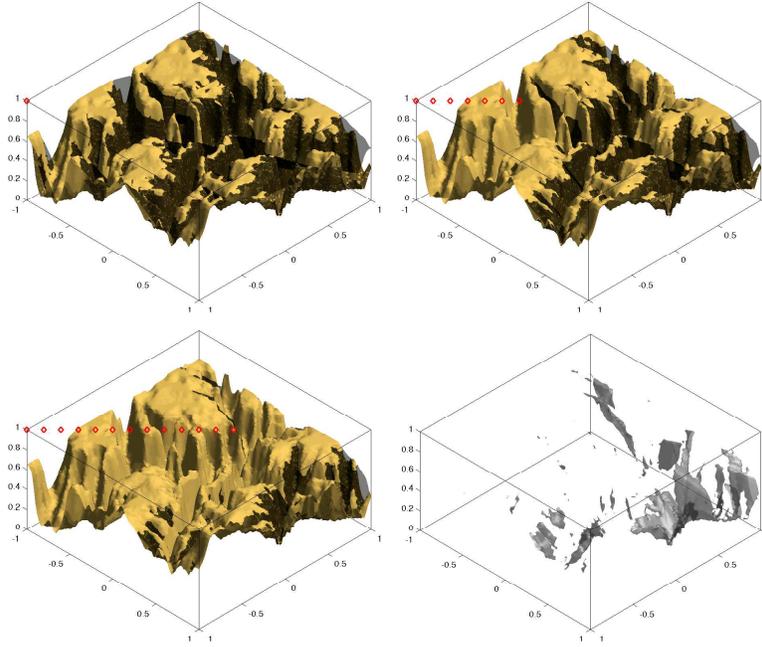


FIGURE 3.6. By taking the intersection of the occlusion during a trajectory of the observer, we can find the cumulative occlusion easily and efficiently. The following pictures show a progression of the cumulative occlusion subject to an observer (“spy plane”) moving across a region of Grand Canyon.

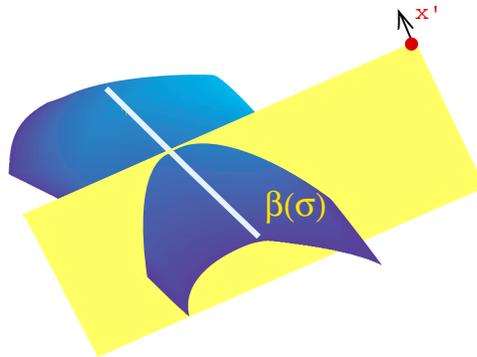


FIGURE 3.7.

where

$$\mathbf{n}^\perp(\mathbf{x}) := \begin{pmatrix} \phi_{x_2}(\mathbf{x}) \\ -\phi_{x_1}(\mathbf{x}) \end{pmatrix} / |\nabla\phi(\mathbf{x})|,$$

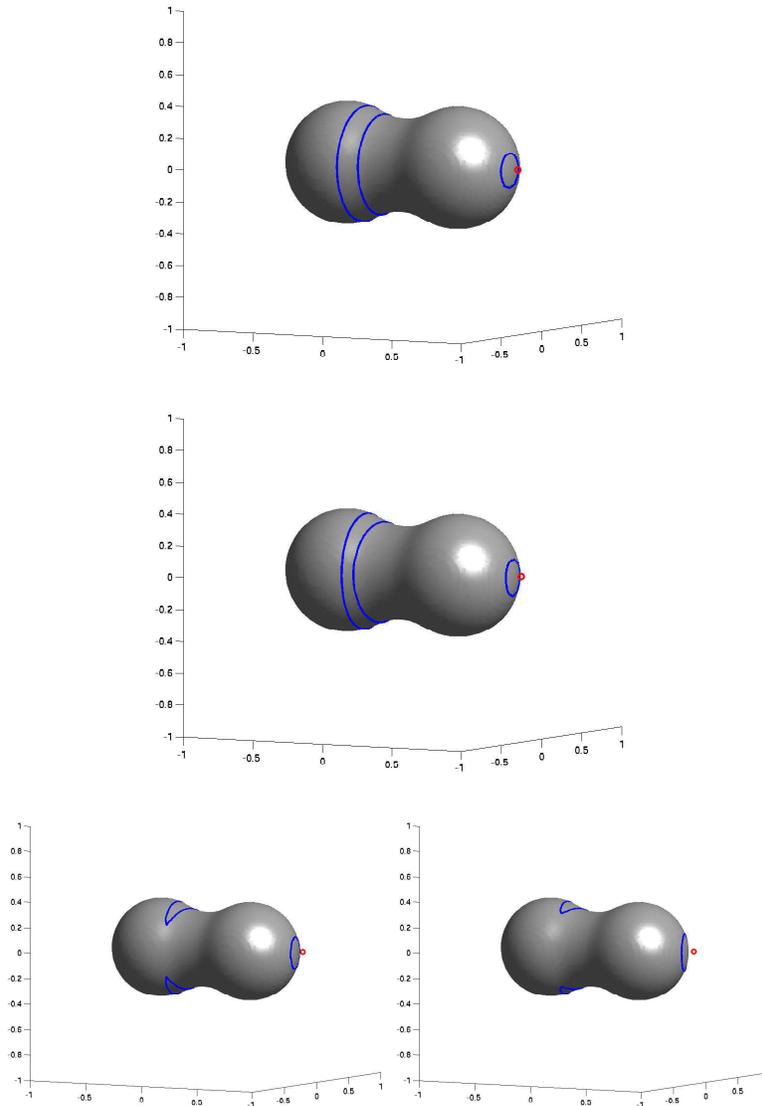


FIGURE 3.8. A example of moving silhouette around a nonconvex occluder. Observe that the silhouette curves break and change topology.

and similarly for v^\perp . See Section 5.5 for a detailed derivation. See Figure 3.9 for a computational result using this formula. We notice that these constraints also tell us how the shadow boundaries should move.

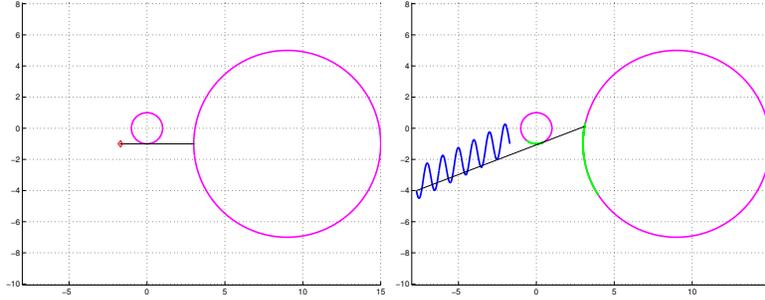


FIGURE 3.9. A result of tracking the terminator and terminator motion using the formulas derived in this paper. The blue curve represents the trajectory of the vantage point and the green curves represent the paths of the terminator and terminator. The black line links the current position of the vantage point and the terminator; it shows that the colinearity of the vantage point, the horizon and its terminator is preserved.

In three dimensions, we can reduce the instantaneous motion to a two dimensional problem on the “right” section of the surface following the reasoning given in the previous subsection.

Motions of the shadow boundaries. How does the shadow move in space? We can constrain a point on the shadow boundary to move only normal to the viewing direction (ergo, the shadow boundary):

$$(3.7) \quad \begin{cases} \mathbf{x}' \cdot \mathbf{v} = 0, \\ \frac{\mathbf{x} - \mathbf{x}_o}{|\mathbf{x} - \mathbf{x}_o|} = \frac{\mathbf{x}^* - \mathbf{x}_o}{|\mathbf{x}^* - \mathbf{x}_o|}. \end{cases}$$

Motions of horizons and terminators of dynamic surfaces. We remark that we are able to derive the motion laws of the visible contours even when the occluders are changing shapes. In this case, the embedding level set function ϕ is a function of space and time, $\phi(\mathbf{x}, t)$ and differentiating formulas (3.2) and (3.5) with respect to t will bring $\phi(\mathbf{x}, t)$ into the equations.

3.4. Analysis of the motions. The formulae derived in the previous subsection can be regarded as a Lagrangian description of the horizon/terminator motion. We extend the velocity to the domain near the surfaces and obtain the corresponding velocity field $\mathbf{v}(\mathbf{x})$. We then evolve the level set function(s) u in question by

$$u_t + \mathbf{v} \cdot \nabla u = 0.$$

The velocity fields for a horizon and its terminator do not depend on the function u . In horizon motion, we evolve \tilde{h} , the velocity is a function of position,

time, $\dot{\mathbf{x}}_o$, and the derivatives of ϕ , i.e. $\mathbf{v} = \mathbf{v}(t, \mathbf{x}, \mathbf{x}_o, \nabla\phi, D^2\phi)$. Furthermore, the level set function to be evolved is \tilde{h} . In terminator motion, we evolve $\tilde{\psi}$, we have $\tilde{\mathbf{v}} = \tilde{\mathbf{v}}(t, \mathbf{x}, \mathbf{x}_o, \nabla\phi, D^2\phi, \tilde{h})$. Therefore, we are evolving the following two level set equations:

$$\begin{aligned}\tilde{h}_t + \mathbf{v}(t, \mathbf{x}, \mathbf{x}_o, \nabla\phi, D^2\phi) \cdot \nabla\tilde{h} &= 0, \\ \tilde{\psi}_t + \tilde{\mathbf{v}}(t, \mathbf{x}, \mathbf{x}_o, \nabla\phi, D^2\phi, \tilde{h}) \cdot \nabla\tilde{\psi} &= 0.\end{aligned}$$

These are simple convection equations whose viscosity solutions are well studied, provided that the velocity fields are bounded. We only have to be careful near singularities.

Formula (3.3) reveals a few interesting facts. First, we notice that the speed of the horizon motion is inversely proportional to the normal curvature in the viewing direction and to the distance between the horizon and the vantage point. If the vantage point is moving in the tangent direction \mathbf{v} , the horizon will not move (since $\dot{\mathbf{x}}_o \cdot \mathbf{n} = 0$). The speed of the horizon motion becomes singular if the curvature of the surface at the horizon location becomes zero. On strictly convex objects, this will never happen. If we restrict our analysis to a single connected smooth non-convex object, we see easily that at the instance in which a horizon point moves into the location where $\kappa = 0$, a neighborhood of this location becomes completely visible. This signifies the disappearance of the horizon point. If the course of the vantage point is reversed, we get the genesis of a new horizon point.

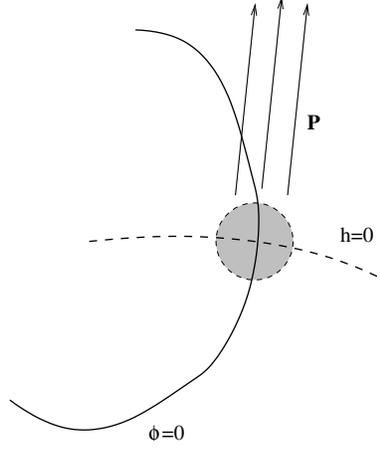
Formula (3.6) tells us that the motion of a terminator point becomes singular when it is a horizon point ($\mathbf{v} \cdot \mathbf{n} = 0$). On a single non-convex smooth surface, this happens precisely when a horizon point and its cast across the concavity collide into each other at the location where $\kappa = 0$. In the setting where there are multiple strictly convex objects, this also describes the changing of the terminator into a visible horizon point which is previously invisible. Therefore, the singularities of the horizons and terminators describe a part of their genesis. A complete genesis of the visible contours includes another part, in which a hidden object suddenly becomes visible. We shall discuss this point in a later subsection.

3.5. Relating horizon and its terminator. To move the terminator, following the notation used in the previous section, we need to find $\mathbf{x}^*(\mathbf{x})$ for each point \mathbf{x} on the terminator. Of course, on the continuum level, \mathbf{x} and $\mathbf{x}^*(\mathbf{x})$ are related by

$$\mathbf{x}^*(\mathbf{x}) := \mathbf{x} - \omega(\mathbf{x})\mathbf{v}(\mathbf{x});$$

$\omega(\mathbf{x})$ can be computed by

$$\omega(\mathbf{x}) = |\mathbf{x} - \mathbf{x}_o| - \rho(\mathbf{v}(\mathbf{x})),$$

FIGURE 3.10. Propagating P .

where v and ρ are defined as previously. In fact, finding ρ or ω is equivalent to solving the visibility problem. Here, we propose an implicit method to find the connection between a horizon point and its terminator in a fashion consistent with our PDE approach; i.e. propagating information along the characteristics of a first order PDE.

In general, let $v(\mathbf{x})$ be the ray vector field. We can propagate the any “seed-ing” horizon point along the ray. We will create a vector field P by the following procedures:

- (1) Set $P = P_0 \notin \Omega$.
- (2) Let \tilde{h} be defined as in (3.1), T_a be a thin tube of radius a around the horizon, and $\partial H_a = T_a \cap \{\tilde{h} = 0\}$. Solve by upwinding

$$\vec{\nabla} P \cdot v(\mathbf{x}) = 0, \quad P(\mathbf{z}) = \mathbf{z}, \forall \mathbf{z} \in \partial H_a \cup \{\mathbf{x}_o\}.$$

Again, from the method of characteristics, P is extended constant to the ray direction from the horizon. See Figure 3.10. The solution P is a piecewise continuous function that is continuous around the terminator. When we move the points \mathbf{x} near terminator, we also move the points $P(\mathbf{x})$, which are points near the horizon. By continuity, we will have the right motion with the terminator.

3.6. Reinitialization and emergence-time estimate. It can be easily seen from figure 3.11 that a completely hidden object may suddenly become visible at a later time during the journey of the vantage point. At the time of emergence, we need to reinitialize our algorithm, i.e., we need to find the visible contours on the newly emerged surfaces to get the correct visibility information. Assuming that we are

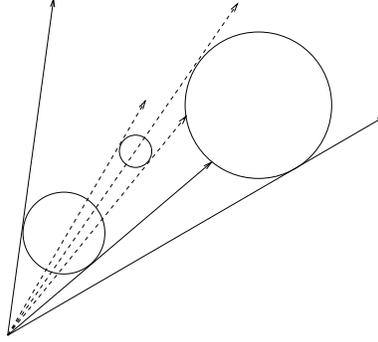


FIGURE 3.11. Model scenario I

merely tracking the visibility boundaries on the objects. How do we know when to initialize? We can formalize the reasoning as follows. We define the map $\mathcal{G}^{-1} : \mathcal{S}^{d-1} \mapsto \{\mathbf{x}_i\}_{i=0}^{N(\theta)}$ such that

$$\mathcal{G}^{-1}(\theta) := \{\mathbf{x}_i \in \mathbb{R}^d : \nabla\phi(\mathbf{x}_i)/|\nabla\phi(\mathbf{x}_i)| = \theta \text{ and } \phi(\mathbf{x}_i) = 0\}.$$

This map is the inverse of the Gauss map in the case that $\{\phi = 0\}$ is a strictly convex hypersurface. Let \mathcal{S} be the set containing x and x^* . We reinitialize whenever there exists an $x \in \mathcal{S}$ such that $\exists y \in \mathcal{G}(v(\mathbf{y}))$ with $\mathbf{x}^*(\mathbf{x}) \prec y \prec \mathbf{x}$. This provides an explicit criterion for reinitialization, but it is certainly not a trivial task. In our implicit formulation, we have more information about the spatial structure of the occluders that we are able to derive an estimate on the time of emergence of a hidden object by using the knowledge of : 1) how the shadow moves; 2) how far a hidden surface is from the shadow boundaries.

Given current vantage point position and its motion, we want to estimate the emergence time for an object that is occluded. We begin by assuming the the curvatures of the surfaces locally around the regions of interest are constant. The diagram in Figure 3.12 shows a model configuration: The small circle is initially occluded by the larger circle on the left. We want to estimate the time interval δt between this instance t_0 and the time $t_1 = t_0 + \delta t$ when the small circle first emerges into the scene.

Following the discussion above, consider a point \mathbf{y} on the shadow boundaries away from the horizon such that $\nabla\phi(\mathbf{y}) \perp v(\mathbf{y})$ and $\nabla\phi(\mathbf{y}) \cdot \nabla\phi(\mathbf{y}(\mathbf{x})) > 0$. This is the point closest to some hidden part of the objects.

For consistency of notation, we will use D in place of \mathbf{y} . Let d be the distance between D and the circle centered at O' . Let ρ and ρ' be the radius of the circle centered at O and O' respectively. Let r denote the distance between D and C' . We

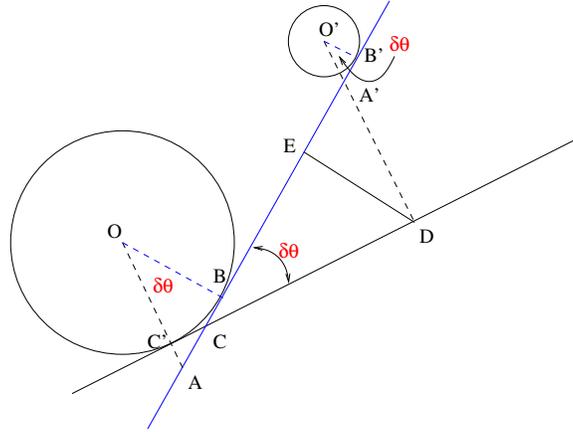


FIGURE 3.12. A diagram for emergence time estimate.

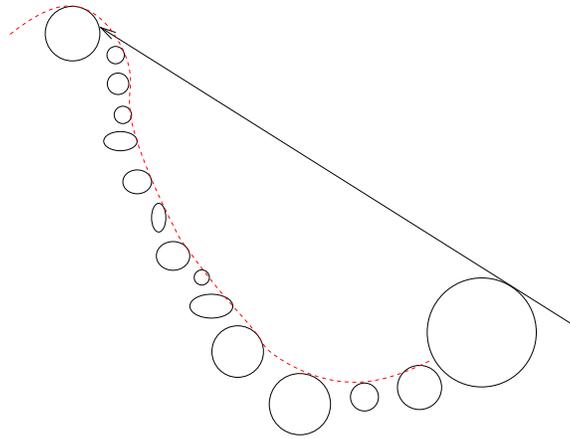


FIGURE 3.13. bad case for moving curves

have the following identities:

$$CC' = \rho \tan \frac{\delta\theta}{2}, \quad CD = r - CC' = r - \rho \tan \frac{\delta\theta}{2}, \quad \implies DE = CD \sin \delta\theta,$$

$$O'A' = \frac{1}{\cos \delta\theta} \rho'.$$

Therefore, we can find $\delta\theta$ from the last two equalities. Since we know how fast the horizon is moving, we can then determine δt .

Let $\dot{\mathbf{x}} = |\dot{\mathbf{x}}|_{\infty}$,

$$\frac{DE}{\delta t} = |\dot{\mathbf{x}}|_{\infty}, \quad \delta t = \frac{DE}{|\dot{\mathbf{x}}|_{\infty}}.$$

3.6.1. *Further considerations — spatial and temporal scales.* We have mentioned in the beginning of this paper that the approach of moving the visible contour may not be more efficient than simply performing implicit ray tracing in general “large world” configurations. Here we construct such a case to validate our arguments.

Consider a nonconvex part of an object as shown by the dashed red curve in Figure 3.13. Suppose the dashed curve is broken down into dense small disconnected components. In the case where the red curve is the nonconvex part of a connected component and with the viewing direction being depicted in Figure 3.13, the terminator will move continuously on the nonconvex part of the object without need for reinitialization.

In essence, this is really an issue of spatial and temporal scales. If the time scale of interest is significantly larger than the distance between these objects, i.e. the vantage point and the occluding surfaces are moving relatively fast when compared to the size of the occluders, then frequent reinitialization is inevitable, and the dynamic approach may be impractical. In this case, we can reconsidered the strategy mentioned earlier of merging these small pieces together, and considered the dynamics in the new “homogenized” setting.

4. CONCLUSION AND FUTURE DIRECTIONS

In this article, we introduced a fast implicit ray tracing algorithm independent of grid geometry and easily parallelizable. This is then extended to a multi-resolution algorithm for near optimal efficiency. Furthermore, we showed that the implicit framework captures accurately the shadow boundaries, which include the horizon and terminator curves. We studied how these objects move when the source point is moving. Explicit formulas which reveal the relations between the motions and the local/global geometry of the given configuration are derived and are tightly coupled with our level set framework for implementation. Also, questions such as “how soon will this hidden object appear” can be answered as a result of our algorithm.

There is a rich pool of applications related to the visibility problem described in this paper. Currently we are working on problems related to navigation, visibility with occluders changing shapes in time, in non-uniform media. Our solutions will combine approaches both from the PDE formulation and the algorithms in computational geometry.

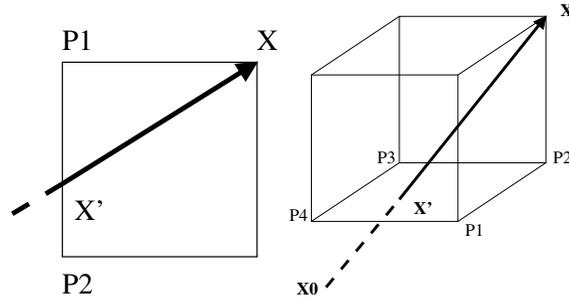


FIGURE 5.1. A demonstration of 2D and 3D interpolation

5. APPENDIX

5.1. Interpolation schemes. Since the majority of visibility applications benefit from the simplicity of Cartesian grids, we need to adapt the algorithm in order to take advantage of this. As described in the algorithm, at each grid point $x = (i, j, k)$, we need to determine an upwind neighbor x' and find the value of $\psi(x')$. In most cases, x' does not lie on the grid. Therefore, we need to interpolate the values of ψ from the grid points closest to x' . For simplicity and speed considerations, we choose to perform linear interpolation in 2D and bilinear interpolation in 3D. In Figure 5.1, we use $\psi(P_1)$ and $\psi(P_2)$ for linear interpolation in the 2D case and use $\psi(P_i)$, $i = 1, 2, 3, 4$, for bilinear interpolation of ψ .

We note that a fast marching or fast sweeping strategy for determining distance from the source point and passing values can be used in place of this interpolation.

Let ψ_{int} be the interpolant near x' , we know that $\psi_{\text{int}}(x') = \psi(x') + O(h^2)$. Thus, the discrete visibility equation (2.3) is in effect

$$\psi(x) = \min(\psi_{\text{int}}(x'), \phi(x)).$$

5.2. Examples of star-shaped updating sequence (sweeping). There are many different ways of implementing a star-shaped updating sequence. One approach is to use the algorithm based on the heap sort strategy [27] to find grid nodes for update based on their distance to the vantage point. However, due the complexity involved with heap sort, this algorithm is not optimal.

Alternatively, we use a sweeping approach in our simulation. For example, let us consider a Cartesian grid in 2D and assume that the vantage point lies on a grid node; we can then consider separately the visibility problem in each of the four quadrants centered at the vantage point. For simplicity, let us assume that the vantage point is at the origin and the grid is represented by the lattice $[-n_x, n_x] \times [-n_y, n_y] \subset \mathbb{Z}^2$. A compact way of writing this sweeping sequence in C/C++ is:

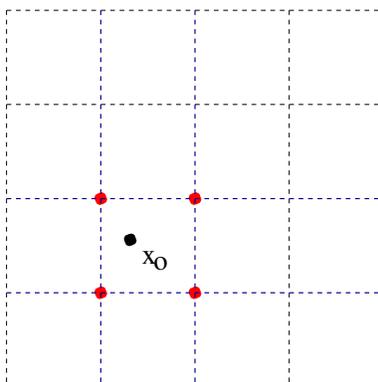


FIGURE 5.2. The red points denote the cell vertices.

```

for (s1=-1; s1<=1; s1+=2)
for (s2=-1; s2<=1; s2+=2)
for (i=0; (s1<0?i>=-nx:i<=nx); i+=s1)
for (j=0; (s2<0?j>=-ny:j<=ny); j+=s2)
  update  $\Psi_{i,j}$ .

```

In the case where \mathbf{x}_0 does not lie on a grid node, we describe an easy modification to the updating sequence above. Let $\mathbf{x}_0 \in I_o := [x_{i_0}, x_{i_0+1}) \times [y_{j_0}, y_{j_0+1})$. Update the values of ψ on the vertices of I_o . Then update the grid nodes in the strips $\{(x_i, y_j) : i = i_0, i_0 + 1 \text{ and } j = -n_y \text{ to } n_y\}$ and $\{(x_i, y_j) : i = -n_x, n_x \text{ and } j = j_0, j_0 + 1\}$. Finally, update the remaining four quadrants independently. See Figure 5.2 for a depiction of this approach.

Finally, we remark that the for loops presented above are meant to demonstrate one possible upwind update sequence for the construction of the solution. In real implementations, one should break it up for better efficiency.

5.3. Finding the curvature of a specified direction. As we argued in Section 3.2, the three dimensional problem of determining the motion of the horizon can be reduced to an instantaneous two dimensional problem. In order to move the horizon in this manner, we need to evaluate the curvature of the surface in the specified direction. Here we present a way to do that.

Let τ be the tangent vector being specified. We want to find the curvature on $\partial\Omega$ in this direction. First let $p(\mathbf{x}; \tau)$ be the plane passing through \mathbf{x} , spanned by $\mathbf{n}(x)$ and τ , and let P be the level set function that embeds this plane. Then

$$\tilde{\tau} = \frac{\nabla\phi \times \nabla P}{|\nabla\phi \times \nabla P|},$$

where $\tilde{\tau}(\mathbf{x}) = \tau$, and the curvature is

$$k_{\tau} \mathbf{n} = \vec{\nabla} \tilde{\tau} \cdot \tilde{\tau}.$$

5.4. Derivation of the dynamics of horizon. We follow the constraints (3.2):

$$\begin{cases} \phi(\mathbf{x}) = 0 \\ (\mathbf{x} - \mathbf{x}_0) \cdot \nabla \phi(\mathbf{x}) = 0 \end{cases}$$

and differentiate with respect to t , we have

$$(5.1) \quad \nabla \phi(\mathbf{x}) \cdot \dot{\mathbf{x}} = 0$$

$$(5.2) \quad (\mathbf{x} - \mathbf{x}_0) \cdot D^2 \phi(\mathbf{x}) \dot{\mathbf{x}} = \dot{\mathbf{x}}_0 \cdot \nabla \phi(\mathbf{x})$$

In 2 space dimensions, these two relations uniquely determine the motion of x with given initial conditions. Writing $(\mathbf{x} - \mathbf{x}_0) = |\mathbf{x} - \mathbf{x}_0| n^\perp(\mathbf{x}) = |\mathbf{x} - \mathbf{x}_0| (-\phi_y, \phi_x) / |\nabla \phi|$, we have

$$\begin{aligned} & \frac{|\mathbf{x} - \mathbf{x}_0|}{|\nabla \phi|} \begin{pmatrix} -\phi_y \\ \phi_x \end{pmatrix} \cdot \begin{pmatrix} \phi_{xx} & \phi_{xy} \\ \phi_{yx} & \phi_{yy} \end{pmatrix} \begin{pmatrix} \dot{x} \\ \dot{y} \end{pmatrix} \\ &= \frac{|\mathbf{x} - \mathbf{x}_0|}{|\nabla \phi|} \begin{pmatrix} -\phi_y \phi_{xx} + \phi_x \phi_{yx} \\ -\phi_y \phi_{xy} + \phi_x \phi_{yy} \end{pmatrix} \cdot \begin{pmatrix} \dot{x} \\ \dot{y} \end{pmatrix}, \end{aligned}$$

and

$$\begin{aligned} & \frac{|\mathbf{x} - \mathbf{x}_0|}{|\nabla \phi|} \begin{pmatrix} \phi_x & \phi_y \\ -\phi_y \phi_{xx} + \phi_x \phi_{yx} & -\phi_y \phi_{xy} + \phi_x \phi_{yy} \end{pmatrix} \begin{pmatrix} \dot{x} \\ \dot{y} \end{pmatrix} \\ &= \begin{pmatrix} 0 \\ \dot{\mathbf{x}}_0 \cdot \nabla \phi(\mathbf{x}) \end{pmatrix} \end{aligned}$$

$$\begin{aligned} \begin{pmatrix} \dot{x} \\ \dot{y} \end{pmatrix} &= \frac{|\nabla \phi|}{|\mathbf{x} - \mathbf{x}_0|} \frac{1}{D} \begin{pmatrix} -\phi_y \phi_{xy} + \phi_x \phi_{yy} & -\phi_y \\ \phi_y \phi_{xx} - \phi_x \phi_{yx} & \phi_x \end{pmatrix} \begin{pmatrix} 0 \\ \dot{\mathbf{x}}_0 \cdot \nabla \phi(\mathbf{x}) \end{pmatrix} \\ &= \frac{|\nabla \phi|}{|\mathbf{x} - \mathbf{x}_0|} \frac{\dot{\mathbf{x}}_0 \cdot \nabla \phi(\mathbf{x})}{D} \begin{pmatrix} -\phi_y \\ \phi_x \end{pmatrix}, \end{aligned}$$

where

$$\begin{aligned} D &= \det \begin{pmatrix} \phi_x & \phi_y \\ -\phi_y \phi_{xx} + \phi_x \phi_{yx} & -\phi_y \phi_{xy} + \phi_x \phi_{yy} \end{pmatrix} \\ &= -\phi_x \phi_y \phi_{xy} + \phi_x^2 \phi_{yy} + \phi_y^2 \phi_{xx} - \phi_x \phi_y \phi_{xy} \\ &= \phi_x^2 \phi_{yy} + \phi_y^2 \phi_{xx} - 2\phi_x \phi_y \phi_{xy}. \end{aligned}$$

Since the curvature of $\partial\Omega$ at x is

$$\begin{aligned}\kappa &= \nabla \cdot \frac{\nabla\phi}{|\nabla\phi|} \\ &= \frac{1}{|\nabla\phi|^3} (\phi_x^2 \phi_{yy} + \phi_y^2 \phi_{xx} - 2\phi_x \phi_y \phi_{xy}),\end{aligned}$$

the motion of x is

$$(5.3) \quad \dot{\mathbf{x}} = \begin{pmatrix} \dot{x} \\ \dot{y} \end{pmatrix} = \frac{1}{\kappa} \frac{\dot{x}_o \cdot n(\mathbf{x})}{|\mathbf{x} - \mathbf{x}_o|} n^\perp(\mathbf{x}).$$

We define $n^\perp(\mathbf{x}) = (\mathbf{x} - \mathbf{x}_o)/|\mathbf{x} - \mathbf{x}_o|$.

Alternatively, we can write \dot{x} in a slightly different form:

$$\dot{x} = \frac{P_{\nabla\phi} \nabla^2 \phi(x - x_0)}{|P_{\nabla\phi} \nabla^2 \phi(x - x_0)|^2} (\dot{x}_o \cdot \nabla\phi),$$

where P_v is the orthogonal projection matrix projecting vectors to the plane with normal vector v . Let us check this expression for the velocity of the curve. Note $\nabla\phi \cdot \dot{x} = 0$ since $v \cdot P_v w = 0$ for all vectors v and w . Also, $\nabla^2 \phi(x - x_0) \cdot \dot{x} = \nabla\phi \cdot \dot{x}_o$ is satisfied since $P_v w \cdot w = |P_v w|^2$ for all vectors v and w . Thus this velocity is valid and is the first form in our alternate derivation.

Geometric interpretation. If x indeed represents the position of the curve, then $x - x_0$ is tangent to the object surface at x and so $x - x_0 = P_{\nabla\phi}(x - x_0)$. Making this replacement above gives our second form for the velocity,

$$\dot{x} = \frac{P_{\nabla\phi} \nabla^2 \phi P_{\nabla\phi}(x - x_0)}{|P_{\nabla\phi} \nabla^2 \phi P_{\nabla\phi}(x - x_0)|^2} (\dot{x}_o \cdot \nabla\phi).$$

This form is particularly nice because we know the second fundamental form in a level set framework is transformed to

$$II = \frac{1}{|\nabla\phi|} P_{\nabla\phi} \nabla^2 \phi P_{\nabla\phi}.$$

Thus we can rewrite the velocity in its final form,

$$\dot{x} = \frac{II(x - x_0)}{|II(x - x_0)|^2} \left(\dot{x}_o \cdot \frac{\nabla\phi}{|\nabla\phi|} \right).$$

IIv evaluated at a point represents the change in the normals of the object surface in the direction of v at that point.

5.5. Derivation of the dynamics of the terminator . We assume that the level sets of ϕ near \mathbf{x} are smooth curves and are not tangent to v . In two space dimension, we have two equations that determine the dynamics of \mathbf{x} :

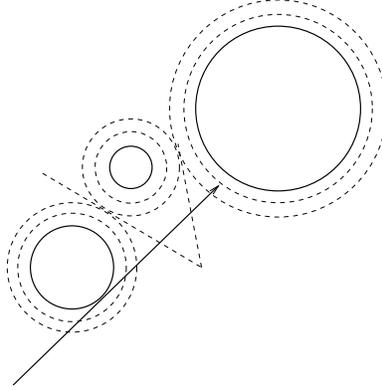


FIGURE 5.3. Model terminator scenario

$$(5.4) \quad \begin{cases} \phi(\mathbf{x}) = \text{constant} \\ \mathbf{v} = \tilde{\mathbf{v}} \end{cases}$$

Let \mathbf{r} and $\tilde{\mathbf{r}}$ denote $(\mathbf{x} - \mathbf{x}_0)$ and $(\tilde{\mathbf{x}} - \mathbf{x}_0)$ respectively. Differentiating these equations, we arrive at:

$$\begin{aligned} \nabla\phi(\mathbf{x}) \cdot \mathbf{x}' &= 0, \\ \frac{\mathbf{r}'}{|\mathbf{r}|} - \frac{\mathbf{r}}{|\mathbf{r}|^2} \mathbf{v} \cdot \mathbf{r}' &= \frac{\tilde{\mathbf{r}}'}{|\tilde{\mathbf{r}}|} - \frac{\tilde{\mathbf{r}}}{|\tilde{\mathbf{r}}|^2} \tilde{\mathbf{v}} \cdot \tilde{\mathbf{r}}'. \end{aligned}$$

Notice that the term

$$\mathbf{v} \cdot \mathbf{r}' \frac{\mathbf{r}}{|\mathbf{r}|} = (\mathbf{v} \cdot \mathbf{r}') \mathbf{v}$$

is \mathbf{r}' projected onto the unit vector \mathbf{v} . Therefore, the left hand side denotes the projection of \mathbf{r}' onto the unit vector \mathbf{v}^\perp :

$$\frac{1}{|\mathbf{r}|} (\mathbf{r}' - \mathbf{v} \cdot \mathbf{r}' \mathbf{v}) = \frac{\mathbf{r}' \cdot \mathbf{v}^\perp}{|\mathbf{r}|} = \frac{1}{|\mathbf{r}|} P_{\mathbf{v}} \mathbf{r}'.$$

Similarly, with the right hand side, we have the equation:

$$\frac{1}{|\tilde{\mathbf{r}}|} P_{\tilde{\mathbf{v}}} \tilde{\mathbf{r}}' = \frac{1}{|\tilde{\mathbf{r}}|} P_{\tilde{\mathbf{v}}} \tilde{\mathbf{r}}'.$$

Keeping in mind that we want to solve for \mathbf{x}' , we move every other term to the right hand side and arrive at

$$\begin{aligned} P_{\mathbf{v}} \mathbf{x}' &= \frac{|\mathbf{r}|}{|\tilde{\mathbf{r}}|} P_{\tilde{\mathbf{v}}} \tilde{\mathbf{r}}' + P_{\mathbf{v}} \mathbf{x}'_0, \\ \nabla\phi(\mathbf{x}) \cdot \mathbf{x}' &= 0. \end{aligned}$$

In two dimensions, $P_V w = (w \cdot v^\perp)v^\perp$, and $P_V w \cdot v^\perp = w \cdot v^\perp$, therefore, we have

$$\begin{pmatrix} v_2 & -v_1 \\ \phi_{x_1} & \phi_{x_2} \end{pmatrix} \begin{pmatrix} x'_1 \\ x'_2 \end{pmatrix} = \begin{pmatrix} \frac{|\mathbf{r}|}{|\tilde{\mathbf{r}}|} \tilde{\mathbf{r}}' \cdot v^\perp + \mathbf{x}'_0 \cdot v^\perp \\ 0 \end{pmatrix},$$

and consequently,

$$\begin{aligned} \begin{pmatrix} x'_1 \\ x'_2 \end{pmatrix} &= \frac{1}{\mathbf{v} \cdot \nabla \phi(\mathbf{x})} \cdot \\ &\begin{pmatrix} \phi_{x_2} & v_1 \\ -\phi_{x_1} & v_2 \end{pmatrix} \begin{pmatrix} \frac{|\mathbf{r}|}{|\tilde{\mathbf{r}}|} \tilde{\mathbf{r}}' \cdot v^\perp + \mathbf{x}'_0 \cdot v^\perp \\ 0 \end{pmatrix} \\ &= \frac{1}{\mathbf{v} \cdot \nabla \phi(\mathbf{x})} \\ &\left(\frac{|\mathbf{r}|}{|\tilde{\mathbf{r}}|} \tilde{\mathbf{r}}' \cdot v^\perp + \mathbf{x}'_0 \cdot v^\perp \right) \begin{pmatrix} \phi_{x_2} \\ -\phi_{x_1} \end{pmatrix} \\ &= \frac{\mathbf{n}^\perp(\mathbf{x})}{\mathbf{v} \cdot \mathbf{n}(\mathbf{x})} \left(\frac{|\mathbf{r}|}{|\tilde{\mathbf{r}}|} \tilde{\mathbf{r}}' + \mathbf{x}'_0 \right) \cdot v^\perp \\ &= \frac{\left(\frac{|\mathbf{r}|}{|\tilde{\mathbf{r}}|} \tilde{\mathbf{r}}' + \mathbf{x}'_0 \right) \cdot v^\perp}{\mathbf{v} \cdot \mathbf{n}(\mathbf{x})} \mathbf{n}^\perp(\mathbf{x}) \end{aligned}$$

where

$$\nabla^\perp \phi(\mathbf{x}) := \begin{pmatrix} \phi_{x_2} \\ -\phi_{x_1} \end{pmatrix}, \text{ and } n^\perp := \frac{\nabla^\perp \phi(\mathbf{x})}{|\nabla^\perp \phi(\mathbf{x})|}.$$

5.6. Numerics. We computed the quantities describe in this paper using standard level set technologies. Please refer to [17, 18, 19] for details.

5.7. A list of level set functions used in this paper. We provide a comprehensive list of the level set functions we construct in this paper:

- ϕ : embeds the objects
- $\tilde{h}(\mathbf{x}) := (\mathbf{x} - \mathbf{x}_0) \cdot \nabla \phi(\mathbf{x})$: characterizes the horizon
- $\tilde{\phi} := \max(\phi, -\tilde{h})$: $\{\tilde{\phi} \leq 0\} = \{\phi \leq 0\} \setminus \{\tilde{h} < 0\}$, defines the same visibility as ϕ
- ψ : the visibility map resulting from the implicit ray tracing on ϕ
- $\tilde{\psi}$: the visibility map resulting from the implicit ray tracing on $\tilde{\phi}$, characterizes the terminator
- $P: \mathbb{R}^d \mapsto \mathbb{R}^d$: links horizon to its cast implicitly

REFERENCES

- [1] D. Adalsteinsson and J.A. Sethian. An overview of level set methods for etching, deposition, and lithography development. *IEEE Transactions on Semiconductor Devices*, 10(1), February 1997.

- [2] Pankaj K. Agarwal and Micha Sharir. Ray shooting amidst convex polygons in 2D. *J. Algorithms*, 21(3):508–519, 1996.
- [3] Pankaj K. Agarwal and Micha Sharir. Ray shooting amidst convex polyhedra and polyhedral terrains in three dimensions. *SIAM J. Comput.*, 25(1):100–116, 1996.
- [4] Luis Alvarez, Frédéric Guichard, Pierre-Louis Lions, and Jean-Michel Morel. Axioms and fundamental equations of image processing. *Arch. Rational Mech. Anal.*, 123(3):199–257, 1993.
- [5] M. D. Betterton. Theory of structure formation in snowfields motivated by penitentes, suncups, and dirt cones. *Physical Review E*, 63, 2001.
- [6] Paul Burchard, Li-Tien Cheng, Barry Merriman, and Stanley Osher. Motion of curves in three spatial dimensions using a level set approach. *J. Comput. Phys.*, 170:720–741, 2001.
- [7] Li-Tien Cheng, Paul Burchard, Barry Merriman, and Stanley Osher. Motion of curves constrained on surfaces using a level set approach. *J. Comput. Phys.*, 175:604–644, 2002.
- [8] Roberto Cipolla and Peter Giblin. *Visual motion of curves and surfaces*. Cambridge University Press, 2000.
- [9] Satyan Coorg and Seth Teller. Real-time occlusion culling for models with large occluders. In *ACM Symposium on Interactive 3D Graphics*, 1997.
- [10] Fredo Durand. 3d visibility: Analysis study and applications. *PhD Thesis, MIT*, 1999.
- [11] Fredo Durand, George Drettakis, Joëlle Thollot, and Claude Puech. Conservative visibility preprocessing using extended projections. In *SIGGRAPH*, 2000.
- [12] Ami Harten. ENO schemes with subcell resolution. *J. Comput. Phys.*, 83(1):148–184, 1989.
- [13] Aaron Hertzmann and Denis Zorin. Illustrating smooth surfaces. In *SIGGRAPH*, 2001.
- [14] Hailin Jin, Anthony Yezzi, Yen-Hsi Tsai, Li Tien Cheng, and Stefano Soatto. Estimation of 3d surface shape and smooth radiance from 2d images: a level set approach. *To appear, Journal of Scientific Computing*.
- [15] Stanley Osher, Li-Tien Cheng, Myungjoo Kang, Hyeseon Shim, and Yen-Hsi Tsai. Geometric optics in a phase-space-based level set and Eulerian framework. *J. Comput. Phys.*, 179(2):622–648, 2002.
- [16] Stanley Osher and Ronald Fedkiw. *Level Set Methods and Dynamic Implicit Surfaces*. Springer-Verlag, New York, 2002.
- [17] Stanley Osher and Ronald P. Fedkiw. Level set methods: an overview and some recent results. *J. Comput. Phys.*, 169(2):463–502, 2001.
- [18] Stanley Osher and James A. Sethian. Fronts propagating with curvature-dependent speed: algorithms based on Hamilton-Jacobi formulations. *J. Comput. Phys.*, 79(1):12–49, 1988.
- [19] Stanley Osher and Chi-Wang Shu. High-order essentially nonoscillatory schemes for Hamilton-Jacobi equations. *SIAM J. Numer. Anal.*, 28(4):907–922, 1991.
- [20] G. Sapiro, R. Kimmel, D. Shaked, B. B. Kimia, and A. M. Bruckstein. Implementing continuous scale morphology via curve evolution. *Pattern Recognition*, 26(9):1363–1372, 1993.
- [21] Gernot Schafler, Julie Dorsey, Xavier Decoret, and Francois X. Sillion. Conservative volumetric visibility with occluder fusion. In *SIGGRAPH*, 2000.
- [22] J. A. Sethian. *Level set methods and fast marching methods*. Cambridge University Press, Cambridge, second edition, 1999. Evolving interfaces in computational geometry, fluid mechanics, computer vision, and materials science.
- [23] Richard Tsai and Li Tien Cheng. Visibility interpolation of point clouds and eno surface reconstruction. *In Preparation*.

- [24] Richard Tsai, Li Tien Cheng, and Stanley Osher. A level set framework for visibility related variational problems. *In Preparation*.
- [25] Yen-Hsi Richard Tsai, Li-Tien Cheng, Paul Burchard, Stanley Osher, and Guillermo Sapiro. Dynamic visibility in an implicit framework. *UCLA CAM Report*, 02(06), 2002.
- [26] Yen-Hsi Richard Tsai, Li-Tien Cheng, Stanley Osher, and Hong-Kai Zhao. Fast sweeping methods for a class of Hamilton-Jacobi equations. *To appear, SIAM J. Numer. Anal.*, 2003.
- [27] John Tsitsiklis. Efficient algorithms for globally optimal trajectories. *IEEE Transactions on Automatic Control*, 40(9):1528–1538, 1995.

RICHARD TSAI, DEPARTMENT OF MATHEMATICS AND PACM, PRINCETON UNIVERSITY,
PRINCETON NJ 08544

E-mail address: ytsai@math.princeton.edu

LI-TIEN CHENG, DEPARTMENT OF MATHEMATICS, UCSD, LA JOLLA, CA 92093-0112.

E-mail address: lcheng@math.ucsd.edu

STANLEY OSHER, DEPARTMENT OF MATHEMATICS, UCLA, LOS ANGELES, CA 90095-
1555

E-mail address: sjo@math.ucla.edu

GUILLERMO SAPIRO, DEPARTMENT OF ELECTRICAL AND COMPUTER ENGINEERING, UNI-
VERSITY OF MINNESOTA, MINNEAPOLIS, MN 55455

E-mail address: guille@ece.umn.edu.