

# Fast Ray Tracing Of Implicit Surfaces

Andrei Sherstyuk <sup>†</sup>

Square USA

55 Merchant Street, Suite 3100, Honolulu HI 96813, USA

Email: andrei@squareusa.com

---

## Abstract

*A ray-tracing algorithm is described for rendering implicit surfaces formed with  $C^1$ -continuous bounded functions  $f(x, y, z)$ . This class of functions includes such popular implicit models as blobby molecules, metaballs, soft objects and convolution surfaces. The algorithm employs analytical methods only, which makes it fast, robust, and numerically stable.*

*An earlier version of this work was presented at the 3rd International Workshop on Implicit Surfaces held in Seattle in 1998.*

---

## 1. Introduction

### 1.1. The problem

An implicit surface is defined as  $S = \{\mathbf{r} \mid F(\mathbf{r}) = 0\}$  where  $F$  is a scalar function  $F : R^3 \rightarrow R$ , defined over all points in 3D space analytically, procedurally, or with elements of both. Such functions are often called field functions or implicit functions. An equation

$$F(\mathbf{r}) = 0 \quad (1)$$

is referred to as an implicit surface equation. The generic form of equation (1) makes implicit surfaces one of the most powerful and flexible modeling tools available.

Ray-tracing <sup>4</sup> allows implicit surfaces to be visualized directly from their models as defined by equation (1), without tessellating it into curves or polygons. The basic operation of ray-tracing is finding ray/surface intersections. Representing the ray parametrically as

$$\mathbf{r}(t) = \mathbf{a} + t\mathbf{b} \quad (2)$$

( $\mathbf{a}$  is the ray's origin and  $\mathbf{b}$  is its direction), an implicit surface equation for all points on a given ray becomes

$$F(\mathbf{r}) = f(\mathbf{a}, \mathbf{b}, t) = 0 \quad (3)$$

or simply

$$f(t) = 0 \quad (4)$$

Equations (3) and (4) must be formulated and solved for millions of rays ( $\mathbf{a}, \mathbf{b}$ ) which poses very heavy computational demands on the rendering system. Few functions (1) currently used in implicit modeling yield closed-form solutions of (4), which necessitates the use of numerical methods and makes the problem even more difficult.

### 1.2. Previous work

There is a multitude of algorithms for ray-tracing implicit surfaces developed to date. A proper classification and evaluation of these algorithms deserves a separate work. An overview of most general methods for ray-tracing implicit surfaces is given by Hart <sup>6</sup>; a number of numerical methods are also described by Mitchell <sup>10</sup>.

With respect to the main implicit equation (1), all algorithms demonstrate various degrees of reliability, speed and generality. Very often, these characteristics are mutually exclusive. The following two algorithms, *ray-marching* and *LG-surfaces*, represent, perhaps, the

<sup>†</sup> This paper was written when the author was at the School of Computer Science and Software Engineering, Monash University, Australia.

most extreme approaches in ray-tracing implicit surfaces.

*Ray-marching.* This is a brute-force method that steps along the ray, evaluating the field function  $f(t)$  on each step. The surface is detected when the sign of  $f(t)$  first changes. The ray-marching algorithm treats field functions  $f(t)$  as true ‘black boxes’ and makes no assumptions about their properties. Ray-marching is one of the most general algorithms which can render anything, given enough time. The robustness is achieved by setting the incremental step low, which makes the algorithm extremely time-consuming. Ray-marching was introduced by Tuy and Tuy<sup>15</sup> for direct visualization of medical data. Perlin and Hoffert<sup>12</sup> used ray-marching to render remarkably complex and realistically looking objects, such as fur, fire and eroded metal, modeled with very noisy functions.

*LG-surfaces.* Kalra and Barr<sup>8</sup> developed an algorithm guaranteed to detect the surface, modeled by functions with computable  $L$  and  $G$  parameters, that represent the Lipschitz constants for the function  $f$  and its derivative  $df/dt$  along the ray. A Lipschitz constant  $\lambda$  is defined for a scalar function  $f$  over region  $A$  as

$$|f(\mathbf{x}) - f(\mathbf{y})| < \lambda \|\mathbf{x} - \mathbf{y}\|, \quad \mathbf{x} \in A, \mathbf{y} \in A$$

A Lipschitz constant  $L$ , computed for a function  $f$ , provides a means to find regions of space where  $f$  is guaranteed not to intersect the surface. A Lipschitz constant  $G$ , computed for the directional derivative  $df/dt$  along a given ray, allows the finding of the intervals of monotonicity of  $f(t)$  and, therefore, the isolation of all roots of  $f(t) = 0$  reliably. The roots are then refined with any well-established method such as regula falsi<sup>13</sup>.

On order to be effective, the *LG-surface* algorithm requires run-time computations of  $L$  and  $G$  for different regions of space and intervals along each ray. To obtain the value of  $L$  over a region of space, the gradient  $\nabla f$  is computed and its magnitude is maximized over the region. For  $G$ , the similar computations must be carried out with  $df/dt$ : the second derivative  $d^2f/dt^2$  is computed and maximized over the interval along the ray. For non-algebraic modeling functions  $f$ , these computations may become prohibitively difficult, even if  $L$  and  $G$  are derived in symbolic form.

There are less demanding algorithms that guarantee to find ray/surface intersections without evaluating second derivatives. For the *sphere-tracing* algorithm, described by Hart<sup>7</sup>, even the first derivatives need not be defined. Sphere-tracing avoids the problem of isolating roots and converges on the surface from one side, using a Lipschitz constant to compute the signed

distance to the surface. Sphere-tracing can render a wide class of surfaces, including fractal, rough and creased ones.

Ray-tracing with interval analysis, introduced by Mitchell<sup>10</sup>, requires run-time evaluation of the first derivative  $df/dt$  to isolate the roots.

It is important to note that all algorithms that bound the rate of change of the implicit functions  $f(t)$ , either with a Lipschitz constant<sup>7, 8</sup> or with derivatives  $df/dt$  along the ray<sup>10</sup> work better when these bounds are as tight as possible. Therefore, they must be computed at run-time for each ray individually and for each interval along this ray, which may not be an easy task to accomplish for complex functions  $f$ . The use of global precomputed values will degrade the efficiency and ultimately will turn the root-isolating algorithms<sup>8, 10</sup> into a simple bisection, and the sphere-tracing algorithm into ray-marching.

To summarize: the most general algorithms with the widest application base<sup>12, 15</sup> are very slow and do not guarantee to locate the surface. On the other hand, reliable methods require auxiliary computations that may become too expensive for complex modeling functions  $f$ .

Next, we provide an overview of the algorithm for ray-tracing implicit surfaces that combines generality, reliability and efficiency.

### 1.3. Algorithm preconditions

The ray-tracing algorithm presented in this paper has been designed to render implicit surfaces modeled with the following equation:

$$F(\mathbf{r}) = -T + \sum_{i=1}^N f_i(\mathbf{r}) \quad (5)$$

where  $T$  is the iso-potential value and each constituent function  $f_i(\mathbf{r})$ , when parameterized along an arbitrary ray (2) as  $f(t)$ , satisfies the following conditions:

1.  $f(t)$  is  $C^1$ -continuous for all  $t$ ,
2.  $f(t)$  and  $df(t)/dt$  have non-zero values only over a finite set of non-overlapping intervals  $[t_{1_i}, t_{2_i}]$ ,  $i = 1, \dots, k$ .

The second condition implies that each object represented by its function  $f$  can be enclosed by a bounding volume (of not necessarily finite size). For example, the field function of an infinite implicit plane may be enclosed in a co-planar infinite slab.

In general, the algorithm requires that both  $f(t)$  and  $f'(t)$  be defined. In special cases when  $f$  is symmetric about the midpoint of each interval  $[t_1, t_2]$ , the

derivative  $f'(t)$  is not required. Examples are: implicit points, lines, tori.

The class of modeling functions that meets the said conditions (and, therefore, can be rendered by our algorithm) is very wide and includes such well-known models such as blobby molecules<sup>1</sup>, metaballs<sup>11</sup>, soft objects<sup>16</sup> and convolution surfaces<sup>2</sup>.

#### 1.4. Algorithm postconditions

The implicit surfaces rendered by the algorithm will exhibit:

##### 1. *Smoothness*

The algorithm will render the object as a  $C^1$ -continuous surface. The surface is guaranteed not to contain pixel dropouts and shading will be smooth along the surface.

##### 2. *Fine features*

The algorithm will render fine features of the implicit surface with the same fidelity as for conventional, non-implicit primitives. The algorithm will miss the implicit surface if its modeling function  $f(\mathbf{r})$  has a bounding volume so small that it slips between the sampling rays. This is the general problem of point-sampling methods that is well understood and may be solved using stochastic or oversampling methods<sup>4</sup>.

##### 3. *Limited accuracy*

The rendered surface may slightly deviate from its true location as implied by the modeling equation (5). The error is individual for each field function  $f_i$  and is not cumulative.

The rest of the paper is organized as follows. The next section provides a detailed description of the algorithm. Section 3 provides a discussion of errors. Implementation issues and speed-up techniques are given in Section 4 and illustrated by practical examples in Section 5. Some possible improvements are suggested in Section 6.

## 2. The algorithm

The algorithm is based on the original work by Nishimura et al., who developed a very efficient, though highly specialized, algorithm for ray-tracing *metaballs*<sup>11</sup>. We first describe Nishimura's algorithm and then show how it can be extended to render implicit functions that meet our pre-conditions.

### 2.1. Ray-tracing algorithm for metaballs

In the *metaball* model, the constituents  $f_i$  of the iso-surface equation (5) are point potentials represented

by piecewise quadratics:

$$f(\mathbf{r}) = \begin{cases} 1 - 3\left(\frac{r}{R}\right)^2 & 0 \leq r \leq \frac{R}{3}; \\ \frac{3}{2}\left(1 - \frac{r}{R}\right)^2 & \frac{R}{3} < r \leq R; \\ 0 & r > R; \end{cases} \quad (6)$$

where  $R$  is radius of influence of the metaball and  $r$  is the distance from its center to point  $\mathbf{r}$ .

To find the ray/surface intersections, the ray equation (2) is substituted into the potential function (6), yielding at most three piecewise quadratic polynomials per metaball that describe its field along the ray. When all metaballs have been processed in this manner, the whole extent of the ray inside their collective field becomes sliced into a set of intervals with corresponding polynomials derived from equation (6). The algorithm walks through these intervals, building and solving the iso-surface equation (5). Since all components are represented by quadratic polynomials, the collective iso-surface equation is also a quadratic, and all roots (hence intersections) can be found analytically.

A similar technique is described by Wyvill and Trotman<sup>18</sup>. They modeled point sources by pieces of polynomials of degree 6, and solved the implicit equations for roots using Laguerre's method.

### 2.2. Generalization of the algorithm for metaballs

The nature of the algorithm described above is that it may be applied to solve iso-surface equations generated by *any field function*, provided it can be represented as a sum of polynomials with ray distance  $t$  as argument. For metaballs, this representation is straightforward, because the modeling functions (6) are already specified in polynomial form. For an arbitrary function  $f$ , additional processing may be required in order to provide this representation. This can be achieved using polynomial approximation.

Weierstrass's theorem states that if  $f$  is a continuous function on the interval  $[t_1, t_2]$ , then for any  $\epsilon > 0$  there exists a polynomial  $p$  such that

$$\max|f(t) - p(t)| < \epsilon, \quad t_1 \leq t \leq t_2 \quad (7)$$

which simply means that any continuous function  $f$  may be approximated on a closed interval by a polynomial  $p$  as closely as required.

In our case, the interval  $[t_1, t_2]$  is obtained via intersecting the ray with the bounding volume of the function  $f$  (see Figure 1). Finding the right approximating polynomial is a matter of balancing between the following constraints:

- The degree of the approximating polynomials  $p(t)$  must be at most 4, so that analytical methods<sup>14</sup> can be used to solve for roots  $t$ .
- For piecewise polynomials, the representation must be  $C^1$ -continuous to guarantee a smooth shading of the surface.
- The number of evaluations of the function  $f(t)$  required to compute the polynomial representation, should be as low as possible to keep the overall performance high.

Interpolation using Hermite polynomials provides a feasible solution to this problem. The Hermite interpolant of a function  $f$  is the polynomial  $p$  of least degree for which

$$\begin{aligned} p(x_i) &= f_i := f(x_i), \\ p'(x_i) &= f'_i := f'(x_i), \end{aligned} \quad (8)$$

at given node points  $x_i$ ,  $i = 1, \dots, n$ .


For our purposes, the Hermite interpolants can be computed as follows. The initial interval  $[t_1, t_2]$ , obtained via intersecting the ray with the bounding volume, is divided at the midpoint  $t_{mid}$  (see Figure 1). For endpoints of the first sub-interval  $[t_1, t_{mid}]$  conditions (8) are written as

$$\begin{aligned} f_1 &= 0, \\ f'_1 &= 0, \\ f_{mid} &= f(t_{mid}), \\ f'_{mid} &= f'(t_{mid}) \end{aligned} \quad (9)$$

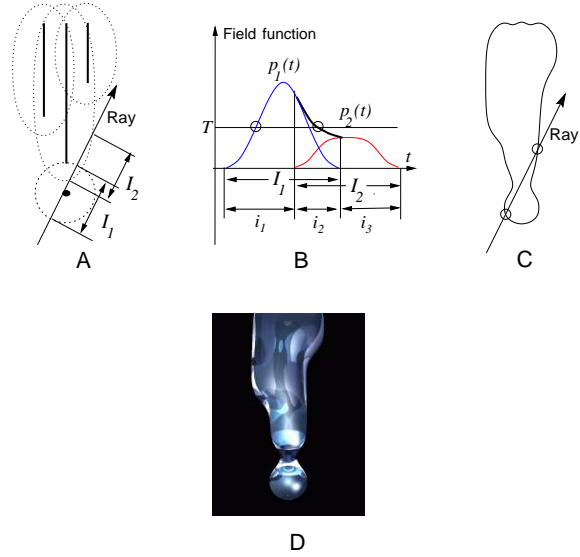
because the function  $f$  is known to have zero values and zero derivatives at the boundaries of its zone of influence. If the function  $f$  is symmetric about the midpoint  $t_{mid}$ , the second derivative condition is reduced to

$$f'_{mid} = 0$$

Equations (9) solved for sub-intervals  $[t_1, t_{mid}]$  and  $[t_{mid}, t_2]$  produce a piecewise representation of the function  $f$  over interval  $[t_1, t_2]$  that satisfy the require-



**Figure 1:** Polynomial approximation over  $[t_1, t_2]$ .



**Figure 2:** Rendering an implicit icicle: (A) intersection with bounding volumes; (B) interpolation of contributing components by polynomials  $p_i(t)$ ; (C) iso-surface at level  $T$ ; (D) shaded image.

ments listed above: it is of low-degree,  $C^1$ -continuous and computationally efficient.

### 2.3. An example

To demonstrate, consider the simple example of finding all intersections between a ray and a surface modeled with three implicit line segments and one point potential (Figure 2 A). We assume that the functions  $f_{point}(t)$  and  $f_{line}(t)$  are defined at any point  $t$  on the ray. The exact expressions for the field function  $f_{point}(t)$  are given by many researchers<sup>1, 16, 11</sup>; implicit functions for a line segment may be found elsewhere<sup>9</sup>.

First, the ray is intersected with the bounding volumes of all modeling functions  $f_i$  (Figure 2A). The resultant intervals  $I_1$  and  $I_2$  define the geometric location along the ray where the field is considered non-zero. Next, for each interval  $I_i$ , the corresponding field function  $f_i$  is interpolated by polynomials  $p_i$  (in this example shown as quartics without loss of generality) (Figure 2B). Finally, all intervals  $I_i$  are intersected and sorted along the ray, yielding a sequence  $i_1, i_2, i_3$  (Figure 2B).

At this point, the algorithm is ready to proceed with the root-finding. The iso-surface equations are built and solved for roots  $t$  in all intervals:

Interval	Field equation	Number of valid roots
$i_1$	$p_1(t) = T$	1
$i_2$	$p_1(t) + p_2(t) = T$	1
$i_3$	$p_1(t) = T$	0

In general, the number of roots per interval may be as high as the highest degree of all interpolating polynomials  $p_i(t)$  defined over this interval. These roots may also occur outside the interval. Therefore, for each interval, the algorithm validates all roots by checking if they belong to the interval. In our example, only valid roots are marked by circles in Figure 2(B and C). The corresponding points give the location of the iso-surface at level  $T$  (Figure 2C). The shaded image is shown in Figure 2D.

## 2.4. Shading and texturing

The normal vector  $\mathbf{n}$  at the ray/surface intersection point  $\mathbf{x}$  is

$$\mathbf{n} = - \sum_{i=1}^N w_i \frac{\nabla f_i(\mathbf{x})}{\|\nabla f_i(\mathbf{x})\|} \quad (10)$$

where the scalar weights, obtained as  $w_i = f_i(\mathbf{x})$ , are normalized to sum up to the threshold value  $T$  (which equals 1 by convention):

$$\sum_{i=1}^N w_i = T \quad (11)$$

Scalar weights  $w_i$  are also used to interpolate between photometric characteristics  $C$  of materials associated with the modeling primitives  $f_i$ :

$$C = \sum_{i=1}^N w_i C_i \quad (12)$$

The choice of parameters  $C$  depends on the lighting model. They usually include ambient, diffuse and specular colors, transparency, reflectivity and other data.

The normalization of  $w_i$  is necessary because of the approximate nature of the algorithm. The actual value of the field at the intersection point  $\mathbf{x}$ , as found by the algorithm, may slightly differ from  $T$ . Thus, the weights  $w_i$  must be scaled to ensure that all photometric parameters and normal vector components blend correctly. ( Obviously the normal vector may be obtained without computing the weights  $w_i$ . In the current implementation of the algorithm, however, it is more convenient to compute the magnitude and direction of the normal vector separately for all components  $f_i$  of the iso-surface.)

Surface texturing, both flat and solid, can also be

performed. Applying solid textures is straightforward and depends on the location of the surface only. Flat textures are more difficult, because local texture coordinates  $(U, V)$  must be computed in the local space of all participating constituents  $f_i$  and then combined together and re-mapped onto the resulting surface. This re-mapping is not easy and is addressed in works by Wyvill et al.<sup>17</sup> Zonenschein et al.<sup>19</sup> and others.

We used the simple fact that each function  $f_i$  is accompanied by its own bounding volume, usually a simple geometric object. Therefore, the  $UV$ -values for the bounding volumes are obtained first and then used to texture the surface parameters  $C$  for each participating  $f_i$  separately. The new textured values  $C_i$  are finally mixed as defined by equation (12).

## 3. Error analysis

As with any technique that involves approximation, it is important to obtain bounds for the errors. In the case of Hermite interpolation on  $n$  nodes  $x_1, x_2, \dots, x_n$  the error is

$$\epsilon(t) = \frac{f^{(2n)}(\xi)}{(2n)!} [L_n(t)]^2 \quad (13)$$

where  $L_n(t) = \prod_{i=1}^n (t - x_i)$  and point  $\xi$  belongs to an interval containing all  $x_i$  and the point of interest  $t$ . Here  $x_i$  denotes the location of an interpolation node and is not to be confused with the endpoints of the initial interval  $[t_1, t_2]$ . The derivation of formula (13) may be found in<sup>3</sup>.

It is apparent from formula (13) that the error may be reduced by increasing the number of node points  $n$  within the initial interval  $[t_1, t_2]$  (see Figure 1). Alternatively, one may fix the number of node points  $n$  and split the initial interval  $[t_1, t_2]$  into smaller sub-intervals, reducing the value of the  $L_n$  term. Both methods reduce the amount of error very efficiently. However, increasing the number of node points  $n$  yields interpolating polynomials of higher degrees  $(2n - 1)$ , which disables the use of analytical root-solvers for  $n > 2$ .


Thus, we have chosen to stay with cubic interpolants ( $n = 2$ ) and increase the number of sub-intervals instead. On each sub-interval  $[x_1, x_2]$ , the error is

$$\epsilon(t) = \frac{f^{(4)}(\xi)}{24} [(t - x_1)(t - x_2)]^2 \quad (14)$$


Since the error is a quartic function of  $|x_1 - x_2|$ , doubling the number of sub-intervals over the initial extent  $[t_1, t_2]$  decreases the error by the factor of 16.

To see the error-control mechanism in practice, consider a following example. A cross is modeled with two


implicit line segments and rendered several times in different views. In the following figures, the leftmost images are rendered with numerical root-finding techniques, using true function evaluations. The middle images are rendered using analytical root-solver and Hermite interpolation. The rightmost images show the rendering error, measured for each pixel as the distance between surface locations as computed by numerical and analytical methods.



The front view. The error is uniform over the image. Two Hermite interpolants are used per line segment. The error is noticeable but insignificant.



The side view. The error is substantially higher in the central areas, where rays travel along the horizontal line segment, yielding longer extents  $[t_1, t_2]$ . The surface becomes deformed.



The side view, revisited. The error is reduced by using four sub-intervals and four Hermite interpolants per line segment. The left and middle images are practically identical.

## 4. Optimizations

The algorithm may be optimized for a particular rendering task. Here are some examples.

### 4.1. Polynomization on demand

If constructive solid geometry (CSG) is not required, the algorithm does not require computation of *all*

polynomials along the ray – it suffices to prepare the intervals  $i_1, \dots, i_n$  only (see Figure 2B). The corresponding interpolants  $p_i$  will be computed on demand, as the algorithm goes through the list of intervals. The process terminates after the first intersection is found. This simple technique may reduce the number of calls to the interpolation routine significantly. In the example given in Figure 2B, the second field function, defined over  $I_2$ , need not be interpolated.

### 4.2. Fast ray/surface rejection test

Before attempting the intersection test, some preliminary interrogation of the modeling functions may help to speed up the intersection test. For example, if the total sum of the maximum contributions from all constituents  $f_i$  is still less than the iso-potential value  $T$ , the iso-surface equation  $\sum_i^n f_i = T$  will have no roots and there will be no intersections in the whole sequence of intervals  $i_1, \dots, i_n$ . This test is especially effective if the modeling functions  $f_i$  are symmetric about the midpoint of their intervals (which is true for implicit points, lines, tori and planes), where the field reaches its maximum value.

A similar root-exclusion test may be used locally for some intervals  $i_j = [x_1, x_2]$ , where the field function  $f$  is known to be monotonic. If  $T$  is not contained between  $f(x_1)$  and  $f(x_2)$ , there will be no intersection in  $i_j$ , and the algorithm may move on to the next  $i_{j+1}$  interval, skipping the interpolation and root-solving stages. To determine if the function  $f = \sum_i^n f_i$  is monotonic, one must check that all its constituents  $f_i$  are consistently non-decreasing (or non-increasing) over that interval.

### 4.3. Volatile and permanent clusters

The modeling functions  $f_i$  that require blending are normally organized in linked lists, or *fusion clusters*, using the terminology of Nishimura et al.<sup>11</sup>. The way these clusters are created may influence the efficiency of the algorithm. The following two types of clusters have been implemented and compared: permanent and volatile.

The first method involves a pre-processing stage, during which all the field functions  $f_i$  in the database are arranged into permanent lists. To create these lists, a connectivity graph is built whose vertices represent all modeling functions  $f_i$  and whose edges are set between objects with intersecting bounding boxes. A depth-first search for fully connected components in this graph extracts all permanent clusters, including degenerate ones for isolated functions. These clusters are permanent because they remain unchanged during the ray-tracing of the whole image. This method

is quite suitable for rendering still pictures and animation sequences where modeling functions  $f_i$  do not change parameters that might affect their zones of influence, and hence the connectivity graph. In the general case, clusters must be decomposed into their constituents and rebuilt again for every new frame. Use of permanent clusters gives better rendering times for simple models, such as in Figure 2.

Alternatively, fusion clusters may be created ‘on the fly’ for each ray. When the ray encounters an implicit function (i.e., intersects its bounding volume), the function is linked to a temporary list. When the traversing of the whole database finishes, this temporary cluster contains all functions  $f_i$  that may contribute to the field along the ray. The cluster is tested for intersections and then destroyed.

Practice shows that volatile clusters are best for rendering complex scenes. Volatile clusters are even more efficient when dynamic memory allocations are replaced by the use of static pools. It is important to note that, regardless of the type of clusters used, the implicit equations formulated for each ray are the same. Choosing between volatile and permanent clusters only changes the way the memory is organized and addressed.

#### 4.4. Reusing interpolants

When shading an intersection point, the algorithm re-evaluates all contributing functions  $f_i$  at that point to compute the weights  $w_i$  (equations (10) and (12)). The speed may be improved significantly by re-using the polynomial representations for each  $f_i$  that were obtained during the ray/surface intersection test. Hermite interpolants  $p_i(t)$  of degree 3 are evaluated at a flat rate of 3 multiplications and 3 additions; true values of field functions  $f_i(t)$  may cost as much as dozens of floating point operations and may contain calls to special functions too.

The following table provides the actual timing results for the model of a coral tree (Figure 6, right), rendered with various optimizations:

Optimization method	Time (min:sec)	Speed-up (%)
None	34:37	—
Interpolation on demand	33:10	4.2
Fast rejection test	31:13	9.8
Static memory pools	31:31	8.9
Reusing interpolants	34:07	1.4
All of the above	28:30	17.7

## 5. Examples

The well-known sphere-flake model of Eric Haines<sup>5</sup>, was re-modeled by replacing the original spheres with Blinn’s blobs

$$f(r) = b e^{-ar^2}$$

where the scalar parameters  $a$  and  $b$  were derived from the radius in isolation for each layer of elements and their blobbiness that controls the blending (see<sup>1</sup> for more detail). To check the speed of our algorithm, each image in Figure 3 was rendered twice: first with the algorithm presented in this paper and then with the *LG*-based algorithm as described by Kalra and Barr<sup>8</sup>. To make a fair comparison, both algorithms were implemented and tested in the identical environment. Special care was taken to implement both methods with the same level of optimization, i.e. obvious things such as multiple function evaluations were carefully avoided. Running times are given in the following table (frame size 512 x 512, supersampling with at most 16 rays per pixel):

Fig	Number of elements	<i>LG</i> -algorithm (min:sec)	Our algorithm (min:sec)
—	1	:57	:46
3A	10	10:27	3:23
3B	91	27:02	7:31
3C	820	50:32	13:53
3D	7381	95:52	31:30

The pseudo-color chart shows the relative amount of time spent rendering the image in Figure 3C with our algorithm (left) and the *LG*-algorithm (right). As expected, silhouette edges do not require extra efforts, because multiple roots at the edges are resolved by analytic root-finding techniques and not by reducing the size of iterative steps as in most numerical methods. For these datasets, our algorithm appeared to be three times faster than the *LG*-algorithm.

It is worth mentioning that a very simple hill-climbing approach renders images 3A - 3D only 1.75 times slower than our algorithm. The hill-climbing method steps along the ray, evaluating pairs  $f(t)$  and  $f(t + \epsilon)$ , until a sign change is detected. Then a standard root-refinement routine (such as regula falsi) is called to close on the root. However, the optimal climbing step  $\epsilon$  had to be found experimentally for each sphere-flake, which makes this method much less attractive.

The next picture demonstrates the small-objects problem. The sea urchin in Figure 5 is modeled by a spherical Gaussian bump and several radial implicit line segments of various lengths. The endpoints of the longest segments are located outside the zone of influ-

ence of the central core, resulting in very sharp spikes that are thinner than the pixel size. The algorithm did not miss these fine features. The right image in Figure 5 gives an example of *UV*-mapping, projected back onto the surface.

The dataset for the images in Figure 6 is based upon the tree model<sup>5</sup>. There are 31 'branches' and 512 'needles' in the dataset. The left image is rendered with conventional primitives in 9 min 36 sec; the right image was re-rendered several times with different optimization techniques as discussed in the previous section, yielding a best time of 28 min 30 sec (frame size 320 x 240, supersampling adaptively). The best ratio of rendering times between the left and the right images is only 2.97, which is very low for an implicit surface of such complexity.

The last image shows the generality of the algorithm. The Hermite crab in Figure 7 is modeled by implicit functions of 5 different types with characteristic sizes a hundred times different. The surface was successfully detected and shaded everywhere, including the fine details, such as the pincer grips.

All images were rendered on a 90MHz Pentium processor.

## 6. Conclusions

The key idea of the ray-tracing algorithm presented in this paper is to keep the modeling implicit equation (5) in polynomial form, so it may be solved quickly during the ray/surface intersection test. The polynomial representation of all implicit functions is obtained via Hermite interpolation.

The algorithm has been implemented as part of the integrated environment for modeling, rendering and animating implicit surfaces *RATS Version 7.30* and has proved very useful. It has been tested with a number of implicit functions, most of which were obtained via a convolution technique described by Bloomenthal and Shoemake in<sup>2</sup> and applied elsewhere<sup>9</sup>. The algorithm has a number of valuable features:

### 1. Modularity

New modeling implicit primitives may easily be added, provided that their functions  $f_i$  meet the general assumptions given in Section 1.3. In short, new implicit primitive =  $f_{new}(\mathbf{r})$  + bounding volume.

### 2. No adjustments required

Changing parameters of the modeling functions does not require re-adjustments of the rendering parameters. This allows the designer to concentrate on the model, not on the rendering.

### 3. Compatibility with other speed-up techniques

Models may consist of a large number of implicit functions  $f_i$ . Since most of them have finite bounds, the rendering speed benefits from space-packing methods, such as grids, in the same manner as for conventional primitives.

### 4. High speed

Rendering speed has always been an issue of great importance, especially when interactive rates are required during designing stage. All implicit objects presented in this paper were created or modified from their original 'explicit' models interactively, using the described algorithm as a viewing tool.

The algorithm may be enhanced by adding an automatic error-control mechanism that would be able to determine if more than two Hermite interpolants are required for each particular field function.

To conclude, we want to emphasize the importance of analytical nature of our algorithm. In its basic version (two Hermite interpolants), the algorithm finds the location of the surface at a flat rate of a *single* field function evaluation plus fixed cost of interpolation and root-solving (Figure 3 shows nearly constant rendering time over the whole image). That compares favorably to numeric methods that require iteration, which may easily become prohibitively expensive for complex non-algebraic modeling functions.


## 7. Acknowledgments

Many thanks to Peter Tischer, Jon McCormack and John Newsome Crossley for critical reading and discussion of the earlier versions of this paper; to anonymous reviewers for suggesting an error analysis of the interpolation scheme. Also, the author is most grateful to all participants of the 3rd International Workshop *Implicit Surfaces '98* held in Seattle this year for the warm welcome of that paper.


## References

1. Blinn J.F., "A Generalization of Algebraic Surface Drawing", *ACM TOG*, Vol. 1, No. 3, July 1982, pp. 235-256.
2. Bloomenthal J. and Shoemake K., "Convolution Surfaces", *SIGGRAPH Proceedings*, Vol. 25, No. 4, July 1991, pp. 251-256.
3. Buchanan J. and Turner P., *Numerical Methods and Analysis*, McGraw-Hill, Inc, 1992.
4. Foley J.D., van Dam A., Feiner S.K. and Hughes J.F., *Computer Graphics, Principles and Practice*, second edition. Reading, Massachusetts: Addison-Wesley, 1990.


5. Haines E., "A Proposal for Standard Graphics Environments", *IEEE Computer Graphics and Applications*, Vol. 7, No. 11, November 1987, pp. 3-5.
6. Hart J.C., "Ray tracing Implicit Surfaces", in *Course Notes 25, SIGGRAPH 1993, Modeling, Visualizing and Animating Implicit Surfaces* pages 13.1-13.15.
7. Hart J.C., "Sphere tracing: A geometric method for the antialiased ray tracing of implicit surfaces", *The Visual Computer*, 12(10), pp. 527-545, December 1996.
8. Kalra D. and Barr A., "Guaranteed Ray Intersection with Implicit Surfaces", *SIGGRAPH Proceedings*, Vol. 23, No. 3, July 1989, pp. 297-306.
9. McCormack J. and Sherstyuk A., "Modeling and Rendering Convolution Surfaces", *Computer Graphics Forum*, Vol.17, No.2, 1998.
10. Mitchell D., "Robust Ray Tracing with Interval Arithmetic", *Proceedings of Graphics Interface '90*, Canadian Information Processing Society, Toronto, 1990, pp. 68-74
11. Nishimura H., Hirai M., Kawai T., Kawata T., Shirakawa I. and Omura K., "Object Modelling by Distribution Function and a Method of Image Generation", *The Transactions of the Institute of Electronics and Communication Engineers of Japan*, 1985, Vol. J68-D, Part 4, pp. 718-725, in Japanese (English translation by Takao Fujiwara available).
12. Perlin K., Hoffert E.M., "Hypertexture", *SIGGRAPH Proceedings*, Vol. 23, July 1989, pp. 253-262.
13. Press W.H., Teukolsky S.A., Vetterling W.T. and Flannery B.P., *Numerical Recipes in C: The Art of Scientific Computing*, second edition. Cambridge University Press, 1992.
14. Schwarze J., "Cubic and Quartic Roots", *Graphics Gems* (editor, Andrew S. Glassner), Academic Press, Cambridge, MA, 1990, pp. 404-407.
15. Tuy H.K. Tuy and Tuy L.T., "Direct 2-D Display of 3-D Objects", *IEEE Computer Graphics and Applications*, Vol. 4, No. 10, October 1984, pp. 29-33.
16. Wyvill G., McPheeters C. and Wyvill B., "Data structure for soft objects", *The Visual Computer*, Vol. 2, pp. 227-234, 1986.
17. Wyvill G., Wyvill B., and McPheeters C. "Solid Texturing of Soft Objects", *IEEE Computer Graphics and Applications*, Vol. 7, No. 12, December 1987, pp. 20-26.
18. Wyvill G. and Trotman A., "Ray Tracing Soft Objects", *CG International 90*, Springer Verlag, pp. 469-475, 1990.
19. Zonenschein R., Gomes J., Velho L., and de Figueiredo L., "Controlling Texture Mapping onto Implicit Surfaces with Particle Systems", *Implicit Surfaces '98 Proceedings*, Seattle, June 1998.




**Figure 3:** Implicit sphere-flakes made with 10, 91, 820 and 7381 Gaussian density functions. The smaller objects have lighter surface colors for better contrast.




**Figure 4:** Time-profiling charts for our algorithm (left) and the LG-algorithm (right). The left image shows almost uniform time complexity over the scene; the right image shows that the rendering was much slower at the silhouette edges. Both images were rendered shooting 1 ray/pixel, resolution 512x512.



**Figure 5:** Two sea-urchins. Left: thin objects are not a problem. Right: UV-mapping, cross-dissolved over blending zones.



**Figure 6:** Coral tree: skeletal model (left) and smoothed implicit surface (right). This model served as a test-bed for various speed optimizations discussed in Section 4. Number of elements 548.



**Figure 7:** Hermite crab, modeled with implicit functions of 5 different types. Notice the use of controlled blending: leg segments do not blend with each other while most other body parts do. 641 elements.