



Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/3761/>

Book Section:

Gamito, M.N. and Maddock, S.C. (2007) A progressive refinement approach for the visualisation of implicit surfaces. In: Braz, J., Ranchordas, A., Araújo, H. and Jorge, J., (eds.) *Advances in Computer Graphics and Computer Vision : International Conferences VISAPP and GRAPP 2006, Setúbal, Portugal, February 25-28, 2006, Revised Selected Papers. Communications in Computer and Information Science, 4 (Part 3)*. Springer, Berlin Heidelberg, pp. 93-108. ISBN: 978-3-540-75274-5.

https://doi.org/10.1007/978-3-540-75274-5_6

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

A PROGRESSIVE REFINEMENT APPROACH FOR THE VISUALISATION OF IMPLICIT SURFACES

Manuel N. Gamito*, Steve C. Maddock
Department of Computer Science, The University of Sheffield
M.Gamito@dcs.shef.ac.uk, S.Maddock@dcs.shef.ac.uk

Keywords: Affine arithmetic, implicit surface, progressive refinement, ray casting.

Abstract: Visualising implicit surfaces with the ray casting method is a slow procedure. The design cycle of a new implicit surface is, therefore, fraught with long latency times as a user must wait for the surface to be rendered before being able to decide what changes should be introduced in the next iteration. In this paper, we present an attempt at reducing the design cycle of an implicit surface modeler by introducing a progressive refinement rendering approach to the visualisation of implicit surfaces. This progressive refinement renderer provides a quick previewing facility. It first displays a low quality estimate of what the final rendering is going to be and, as the computation progresses, increases the quality of this estimate at a steady rate. The progressive refinement algorithm is based on the adaptive subdivision of the viewing frustum into smaller cells. An estimate for the variation of the implicit function inside each cell is obtained with an affine arithmetic range estimation technique. Overall, we show that our progressive refinement approach not only provides the user with visual feedback as the rendering advances but is also capable of completing the image faster than a conventional implicit surface rendering algorithm based on ray casting.

1 INTRODUCTION

Implicit surfaces play an important role in Computer Graphics. Surfaces exhibiting complex topologies, i.e. with many holes or disconnected pieces, can be easily modelled in implicit form. An implicit surface is defined as the set of all points \mathbf{x} that verify the condition $f(\mathbf{x}) = 0$ for some function $f : \mathbb{R}^3 \mapsto \mathbb{R}$. Modelling with implicit surfaces amounts to the construction of an appropriate function f that will generate the desired surface.

Rendering algorithms for implicit surfaces can be broadly divided into meshing algorithms and ray casting algorithms. Meshing algorithms convert an implicit surface to a polygonal mesh format, which can be subsequently rendered in real time with modern graphics processor boards (Lorensen and Cline, 1987; Bloomenthal, 1988; Velho, 1996). Ray casting algorithms compute the projection of an implicit surface on the screen by casting rays from each pixel into three-dimensional space and finding their intersection with the surface (Roth, 1982).

Our ultimate goal is to use implicit surfaces as a tool to model and visualise realistic procedural planets over a very wide range of scales. The function f that generates the surface terrain for such a planet must have fractal scaling properties and exhibit a large amount of small scale detail. Examples of this type of terrain generating function can be found in the Computer Graphics literature (Ebert et al., 2003). In our planet modelling scenario, meshing algorithms are too cumbersome as they generate meshes with a very high polygon count in order to preserve all the visible surface detail. Furthermore, as the viewing distance changes, the amount of surface detail varies accordingly and the whole polygon mesh needs to be regenerated. For these reasons, we have preferred a ray casting approach because of its ability to render the surface directly without the need for an intermediate polygonal representation.

The visualisation of an implicit surface with ray casting is not without its problems, however. When the surface is complex, many iterations have to be performed along each ray in order to locate the intersection point with an acceptable accuracy (Mitchell, 1990). Imaging an implicit surface with ray casting can then become a slow procedure. This is further

*Supported by grant SFRH/BD/16249/2004 from Fundação para a Ciência e a Tecnologia, Portugal.

compounded by the fact that an anti-aliased image requires that many rays be shot for each pixel (Cook, 1989).

We propose to alleviate the long rendering times associated with the modelling and subsequent ray casting of complex fractal surfaces by providing a quick previewer based on a progressive refinement rendering principle. The idea of progressive refinement for image rendering was first formalised in 1986 (Bergman et al., 1986). Progressive refinement rendering has received much attention in the fields of radiosity and global illumination (Cohen et al., 1988; Guo, 1998; Farrugia and Peroche, 2004). Progressive refinement approaches to volume rendering have also been developed (Laur and Hanrahan, 1991; Lippert and Gross, 1995). Our previewer uses progressive rendering to visualise an increasingly better approximation to the final implicit surface. It allows the user to make quick editing decisions without having to wait for a full ray casting solution to be computed. Because the rendering is progressive, the previewer can be terminated as soon as the user is satisfied or not with the look of the surface.

Our progressive refinement previewing method relies on affine arithmetic to compute an estimate of the variation of the implicit function f inside some region (Comba and Stolfi, 1993). Affine arithmetic is a framework for evaluating algebraic functions with arguments that are bounded but otherwise unknown. It is a generalisation of the older interval arithmetic framework (Moore, 1966). Affine arithmetic, when compared against interval arithmetic, is capable of returning much tighter estimates for the variation of a function, given input arguments that vary over the same given range. Affine arithmetic has been used with success in an increasing number of Computer Graphics problems, including the ray casting of implicit surfaces (de Cusatis Jr. et al., 1999). We use a simpler form of affine arithmetic known as *Affine Form 1* (AF1), which we term *reduced affine arithmetic* (Messine, 2002). Reduced affine arithmetic, in the context of ray casting implicit surfaces made from procedural fractal functions, returns the same results as standard affine arithmetic while being faster to compute and requiring smaller data structures.

2 PREVIOUS WORK

One of the best known techniques for previewing implicit surfaces at interactive frame rates is based on the dynamic placement of discs that are tangent to the surface (Witkin and Heckbert, 1994; Hart et al., 2002). The discs are kept apart by the application of repulsive forces and are constrained to remain on the implicit surface. Each disc is also made tangent to

the surface by sharing the surface normal at the point where it is located. This previewing system relies on a characteristic of our visual system whereby we are able to infer the existence of an object based solely on the distribution of a small number of features on the surface of that object. This visual trait only works, however, when the surface of the object is simple and fairly smooth. If the surface is irregular, an apparently random distribution of discs is visible and no object is perceived.

An approximate representation of an implicit surface can be generated by subdividing the space in which the surface is embedded into progressively smaller voxels and using a surface classification technique to identify which voxels are potentially intersecting with the surface. One such spatial subdivision method employs interval arithmetic to perform the surface classification step (Duff, 1992). The subdivision strategy of this method is adapted from an earlier work and is not suitable for interactive previewing (Woodwark and Quinlan, 1982). One must wait for the subdivision to finish before any surface approximation can be visualised unless some additional data processing is added, which will tend to slow down the algorithm. Another spatial subdivision method employs affine arithmetic to perform surface classification and subdivides space with an octree data structure (de Figueiredo and Stolfi, 1996). The octree voxels are rendered from back to front, relative to the viewpoint, with a painter's algorithm. This subdivision strategy is wasteful as it tracks the entire surface through subdivision, including parts that are occluded and that could be safely discarded for a given viewing configuration.

Rather than performing object space subdivision, one can also perform image space subdivision in order to obtain a progressive rendering mechanism. Sample subdivision in image space was originally proposed as an anti-aliasing method for ray tracing (Whitted, 1980). Four rays are shot at the corners of each rectangular sample. If the computed colours for these rays differ by more than some specified amount, the sample is subdivided into four smaller samples and more rays are shot through the corners of the new samples. This type of image space subdivision can also be used for progressive refinement previewing (Painter and Sloan, 1989; Maillot et al., 1992). The problem with image space subdivision algorithms is that they rely entirely on probabilistic methods to determine when to subdivide the image samples. The decision to subdivide a sample is based on a probabilistic analysis of the set of rays traced so far in the neighbourhood of that sample. Because this discrete set of rays is only an approximation of a continuous image intensity distribution, wrong subdivision decisions can sometimes occur.

3 RENDERING WITH PROGRESSIVE REFINEMENT

The main stage of our method consists in the binary subdivision of the space, visible from the camera, into progressively smaller cells that are known to straddle the boundary of the surface. The subdivision mechanism stops as soon as the projected size of a cell on the screen becomes smaller than the size of a pixel. Information about the behaviour of the implicit function f inside a cell is returned by evaluating the function with reduced affine arithmetic. The procedure for rendering implicit surfaces with progressive refinement can be broken down into the following steps:

1. Build an initial cell coincident with the camera's viewing frustum. The near and far clipping planes are determined so as to bound the implicit surface.
2. Recursively subdivide this cell into smaller cells. Discard cells that do not intersect with the implicit surface. Stop subdivision if the size of the cell's projection on the image plane falls below the size of a pixel.
3. Assign the shading value of a cell to all pixels that are contained inside its projection on the image plane. The shading value for a cell is taken from the evaluation of the shading model at the centre point of the cell.

The following sections will explain how each of the steps in our rendering method work, starting with a presentation of the reduced affine arithmetic framework in Section 3.1. We then explain the geometry of a cell inside the camera's viewing frustum (Section 3.2) and how a cell is subdivided and rendered (Sections 3.3 and 3.4). We also explain in Section 3.5 how a region of interest can be optionally defined so as to provide the user with interactive control during the rendering process.

3.1 Reduced Affine Arithmetic

A variable is represented with reduced affine arithmetic (rAA) as a central value plus a series of noise symbols. In contrast to the standard affine arithmetic model, the number of noise symbols is constant and can be used to describe the fundamental degrees of freedom of the problem under consideration (Messine, 2002). In the rendering method that is being described in this paper, the degrees of freedom are the three parameters necessary to locate any point inside the viewing frustum of the camera. These parameters are the horizontal distance u along the image plane, the vertical distance v along the same image plane and the distance t along the ray that passes through the point at (u, v) . A rAA variable \hat{a} has,

therefore, the following representation:

$$\hat{a} = a_0 + a_u e_u + a_v e_v + a_t e_t + a_k e_k. \quad (1)$$

The noise symbols e_u, e_v and e_t are shared between all rAA variables in the system, which allows for the representation of correlation information between rAA variables relative to the u, v and t degrees of freedom. The extra noise symbol e_k is included to account for uncertainties in the \hat{a} variable that are not shared with any other variable.

Operations on rAA variables are performed by updating the a_u, a_v and a_t noise coefficients with their new uncertainties and clumping all other uncertainties into the a_k coefficient. We give an example of how rAA operations work by considering the case of the multiplication between two variables \hat{a} and \hat{b} of the form (1). In the original standard affine arithmetic framework, the result $\hat{c} = \hat{a}\hat{b}$ would be written as:

$$\hat{c} = c_0 + c_u e_u + c_v e_v + c_t e_t + c_{ka} e_{ka} + c_{kb} e_{kb} + c_n e_n. \quad (2)$$

The final error symbols from \hat{a} and \hat{b} were written as e_{ka} and e_{kb} , respectively, to make it clear that they are independent. The new noise symbol e_n is introduced to account for the non-linearity of the multiplication operator. The coefficients for the variable \hat{c} are:

$$\begin{aligned} c_0 &= a_0 b_0, \\ c_u &= a_0 b_u + b_0 a_u, \\ c_v &= a_0 b_v + b_0 a_v, \\ c_t &= a_0 b_t + b_0 a_t, \\ c_{ka} &= b_0 a_k, \\ c_{kb} &= a_0 b_k, \\ c_n &= (|a_u| + |a_v| + |a_t| + |a_k|) \times \\ &\quad (|b_u| + |b_v| + |b_t| + |b_k|). \end{aligned} \quad (3)$$

As a sequence of standard affine arithmetic computations progresses, new noise symbols keep being introduced into the system. For a sufficiently complex expression, the number of noise symbols that have to be considered makes the system increasingly difficult to manage, both in terms of memory requirements and of computational expense. One technique to keep the number of error symbols down to a manageable level is to periodically invoke a procedure called *condensation* (Stolfi and de Figueiredo, 1997). Condensation reduces the number of error symbols of a standard affine arithmetic variable at the cost of destroying correlation information. Reduced AA operations are always followed by a condensation step to remove any extra error symbols that would have been introduced otherwise. Reduced AA can, therefore, be seen as a modification of affine arithmetic that employs an aggressive form of condensation. If the variable \hat{c} in (2)

is condensed into a new variable \hat{d} with only four error symbols, we will have for the coefficients of \hat{d} :

$$\begin{aligned} d_0 &= c_0, \\ d_u &= c_u, \\ d_v &= c_v, \\ d_t &= c_t, \\ d_k &= |c_{ka}| + |c_{kb}| + |c_n|. \end{aligned} \quad (4)$$

The condensed variable \hat{d} is now in the rAA form, according to (1). With the reduced affine arithmetic framework, all operations are always followed by a condensation step to keep a constant number of error symbols for every variable throughout the computation. In practice, all operations in reduced affine arithmetic are modified so that the condensation step (4) is automatically built into them. The multiplication $\hat{c} = \hat{a}\hat{b}$, that in standard affine arithmetic was given by (3), now becomes:

$$\begin{aligned} c_0 &= a_0b_0, \\ c_u &= a_0b_u + b_0a_u, \\ c_v &= a_0b_v + b_0a_v, \\ c_t &= a_0b_t + b_0a_t, \\ c_k &= |a_0b_k| + |b_0a_k| + \\ &\quad (|a_u| + |a_v| + |a_t| + |a_k|) \times \\ &\quad (|b_u| + |b_v| + |b_t| + |b_k|). \end{aligned} \quad (5)$$

Reduced affine arithmetic is more efficient than standard affine arithmetic because it keeps only the required minimum amount of correlation information between all rAA quantities. In our progressive refinement renderer, much faster convergence rates can be obtained towards the final image by using affine arithmetic in reduced form.

For an implicit surface, the value $f(\mathbf{x})$ at some point \mathbf{x} in space can be computed with reduced affine arithmetic. The rAA representation $\hat{\mathbf{x}}$ of the vector \mathbf{x} is a tuple of three rAA coordinates, similar to (1), where each coordinate has its own independent noise symbol e_{k_i} , with $i = 1, 2, 3$. The rAA vector $\hat{\mathbf{x}}$ describes not a point but a region of space spanned by the uncertainties associated with its three coordinates. Evaluation of the expression $\hat{y} = f(\hat{\mathbf{x}})$ leads to a range estimate \hat{y} for the variation of $f(\hat{\mathbf{x}})$ inside the region spanned by $\hat{\mathbf{x}}$. Knowing \hat{y} , the average value \bar{y} and the variance $\langle y \rangle$ for that range estimate can be computed as follows:

$$\bar{y} \triangleq y_0, \quad (6a)$$

$$\langle y \rangle \triangleq |y_u| + |y_v| + |y_t| + |y_k|. \quad (6b)$$

The range estimate \hat{y} is then known to lie inside the interval $[\bar{y} - \langle y \rangle, \bar{y} + \langle y \rangle]$. If this interval contains zero, the region spanned by $\hat{\mathbf{x}}$ may or may not intersect with the implicit function. This is because

affine arithmetic (both in its standard and reduced forms) always computes conservative range estimates and it is possible that the exact range resulting from $f(\hat{\mathbf{x}})$ may be smaller than \hat{y} . What is certain is that if $[\bar{y} - \langle y \rangle, \bar{y} + \langle y \rangle]$ does not contain zero the region spanned by $\hat{\mathbf{x}}$ is either completely inside or completely outside the implicit surface and therefore does not intersect it.

3.2 The Anatomy of a Cell

A cell is a portion of the camera's viewing frustum that results from a recursive subdivision along the u , v and t parameters. Figure 1 depicts the geometry of a cell. It has the shape of a truncated pyramid of quadrangular cross-section, similar to the shape of the viewing frustum itself. Four vectors, taken from the camera's viewing system, are used to define the spatial extent of a cell. These vectors are:

The vector \mathbf{o} This is the location of the camera in the world coordinate system.

The vectors \mathbf{p}_u and \mathbf{p}_v They represent the horizontal and vertical direction along the image plane. The length of these vectors gives the width and height, respectively, of a pixel in the image plane.

The vector \mathbf{p}_t It is the vector from the camera's viewpoint and orthogonal to the image plane. The length of this vector gives the distance from the viewpoint to the image plane.

The vectors \mathbf{p}_u , \mathbf{p}_v and \mathbf{p}_t define a left-handed perspective viewing system. The position of any point \mathbf{x} inside the cell is given by the following inverse perspective transformation:

$$\begin{aligned} \mathbf{x} &= \mathbf{o} + (u\mathbf{p}_u + v\mathbf{p}_v + \mathbf{p}_t)t \\ &= \mathbf{o} + ut\mathbf{p}_u + vt\mathbf{p}_v + t\mathbf{p}_t. \end{aligned} \quad (7)$$




Figure 1: The geometry of a cell. The vectors show the three medial axes of the cell.

The spatial extent of a cell is obtained from the above by having the u , v and t parameters vary over appropriate intervals $[u_a, u_b]$, $[v_a, v_b]$ and $[t_a, t_b]$. We must consider how to compute the rAA representation $\hat{\mathbf{x}}$ of this spatial extent. To do so a change of

variables must first be performed. The rAA variable $\hat{u} = u_0 + u_e e_u$ will span the same interval $[u_a, u_b]$ as u does if we have:

$$u_0 = (u_b + u_a)/2, \quad (8a)$$

$$u_e = (u_b - u_a)/2. \quad (8b)$$

Similar results apply for the v and t parameters. Substituting \hat{u} , \hat{v} and \hat{t} in (7) for u , v and t , we get:

$$\begin{aligned} \mathbf{x} = & \mathbf{o} + t_0 u_0 \mathbf{p}_u + t_0 v_0 \mathbf{p}_v + t_0 \mathbf{p}_t \\ & + t_0 u_e e_u \mathbf{p}_u + t_0 v_e e_v \mathbf{p}_v \\ & + u_0 t_e e_t \mathbf{p}_u + v_0 t_e e_t \mathbf{p}_v + t_e e_t \mathbf{p}_t \\ & + t_e u_e e_u e_t \mathbf{p}_u + t_e v_e e_v e_t \mathbf{p}_v. \end{aligned} \quad (9)$$

The first line of (9) contains only constant terms. The second and third lines contain linear terms of the noise symbols e_u , e_v and e_t . The fourth line contains two non-linear terms $e_u e_t$ and $e_v e_t$, which are a consequence of the non-linearity of the perspective transformation. Since a rAA representation cannot accommodate such non-linear terms they are replaced by the independent noise terms e_{k_1} , e_{k_2} and e_{k_3} for each of the three cartesian coordinates of $\hat{\mathbf{x}}$. The rAA vector $\hat{\mathbf{x}}$ is finally given by:

$$\begin{aligned} \hat{\mathbf{x}} = & \mathbf{o} + t_0(u_0 \mathbf{p}_u + v_0 \mathbf{p}_v + \mathbf{p}_t) \\ & + t_0 u_e \mathbf{p}_u e_u + t_0 v_e \mathbf{p}_v e_v \\ & + t_e(u_0 \mathbf{p}_u + v_0 \mathbf{p}_v + \mathbf{p}_t) e_t \\ & + [x_{k_1} e_{k_1} \quad x_{k_2} e_{k_2} \quad x_{k_3} e_{k_3}]^T, \end{aligned} \quad (10)$$

with

$$x_{k_i} = |t_e u_e p_{u_i}| + |t_e v_e p_{v_i}|, \quad i = 1, 2, 3. \quad (11)$$

A consequence of the non-linearity of the perspective projection and its subsequent approximation with rAA is that the region spanned by $\hat{\mathbf{x}}$ is going to be larger than the spatial extent of the cell. Figure 2 shows the geometry of a cell and the region spanned by its rAA representation in profile. Because the rAA representation has been linearised, its spatial extent is a prism rather than a truncated pyramid. This has further consequences in that the evaluation of $f(\hat{\mathbf{x}})$ is going to include information from the regions of the prism outside the cell and will, therefore, lead to range estimates that are larger than necessary. The linearisation error is more pronounced for cells that exist early in the subdivision process. As subdivision continues and the cells become progressively smaller, their geometry becomes more like that of a prism and the discrepancy with the geometry of $\hat{\mathbf{x}}$ decreases².

The subdivision of a cell proceeds by first choosing one of the three perspective projection parameters

²This can be demonstrated by the fact that the terms $t_e u_e$ and $t_e v_e$ in (9) decrease more rapidly than any of the linear terms u_e , v_e and t_e of the same equation as the latter converge to zero.




Figure 2: The outline of a cell (solid line) and the outline of its rAA representation (dashed line) shown in profile. The rAA representation is a prism that forms a tight enclosure of the cell.

u , v or t and splitting the cell in half along that parameter. This scheme leads to a k - d tree of cells where the sequence of dimensional splits is only determined at run time. The choice of which parameter to split along is based on the average width, height and depth of the cell:

$$\bar{w}_u = 2 t_0 u_e \|\mathbf{p}_u\|, \quad (12a)$$

$$\bar{w}_v = 2 t_0 v_e \|\mathbf{p}_v\|, \quad (12b)$$

$$\bar{w}_t = 2 t_e \|u_0 \mathbf{p}_u + v_0 \mathbf{p}_v + \mathbf{p}_t\|. \quad (12c)$$

If, say, \bar{w}_u is the largest of these three measures, the cell is split along the u parameter. The two child cells will have their u parameters ranging inside the intervals $[u_a, u_0]$ and $[u_0, u_b]$, where $[u_a, u_b]$ was the interval spanned by u in the mother cell. In practice, the factors of 2 in (12) can be ignored without changing the outcome of the subdivision. This subdivision strategy ensures that, after a few iterations, all the cells will have an evenly distributed shape, even when the initial cell is very long and thin.

3.3 The Process of Cell Subdivision

Cell subdivision is implemented in an iterative manner rather than using a recursive procedure. The cells are kept sorted in a priority queue based on their level of subdivision. A cell has priority over another if it has undergone less subdivision. For cells at the same subdivision level, the one that is closer to the camera will have priority. The algorithm starts by placing the initial cell, which corresponds to the complete viewing frustum, on the priority queue. At the start of every new iteration, a cell is removed from the head of the queue. If the extent of the cell's projection on the image plane is larger than the extent of a pixel, the cell is subdivided and its two children are examined. In the opposite case, the cell is considered a leaf cell and is discarded after being rendered. The two conditions that indicate whether a cell should be subdivided are:

with the difference that the secondary queue is now being used. Once this queue becomes empty, the portion of the image inside the ROI is fully rendered and the algorithm returns to subdividing the cells that were left in the primary queue. It is also possible to cancel the ROI at any time by flushing any cells still in the secondary queue back to the primary queue.

3.6 Some Implementation Remarks

The best implementation strategy for our rendering method is to have an application that runs two threads concurrently: a subdivision thread and a rendering thread. An internal image buffer is used to store the rendering of the surface as it is being refined. The subdivision thread requires read-write access to this buffer while the rendering thread requires read-only access to the same buffer. The rendering thread is responsible for periodically updating the graphical output of the application with the latest results from the subdivision thread. Its task is to invoke a single graphics library call that transfers the content of the internal image buffer to the frame buffer of the GPU card. A timer is used to keep a constant frame refresh rate. Except for the periodical invocation of the timer handler routine, the rendering thread remains in a sleep state so that the subdivision thread can use all the CPU resources.

It is possible that, on machines with a small amount of main memory, excessive paging may occur due to the need to store a large number of samples in the priority queue. We have implemented our application on a Pentium 4 1.8GHz with 1Gb of memory. All the results shown in the next section were tested on this computer and it was found that the use of swap memory was never necessary. In any case, it is advisable that the data structure used to hold a sample be as light as possible.

4 RESULTS

Figure 5 on the next page shows four snapshots taken during the progressive refinement rendering of an implicit sphere modulated with a Perlin procedural noise function (Perlin, 2002). The last snapshot shows the final rendering of the surface. The large scale features of the surface become settled quite early and the latter stages of the refinement are mostly concerned with resolving small scale details.

Figure 4 shows an implicit sphere modulated with two and three layers of the Perlin noise function. Table 1 shows the total number of iterations and the computation time for the surfaces that were rendered in Figures 4 and 5. The table also shows the computation time for ray casting the same surfaces by shoot-




Figure 4: An implicit surface with two layers (left) and three layers (right) of a Perlin noise function.

ing a single ray through the centre of each pixel. The number of iterations required to complete the progressive rendering algorithm is largely independent of the complexity of each surface. It depends only on the image resolution and on the percentage of the image that is covered by the projected surface.

Table 1: Rendering statistics for an implicit sphere with several layers of Perlin noise.

<i>Layers</i>	<i>Iterations</i>	<i>Time</i>	<i>Raycasting</i>
1	350759	27.8s	1m10.4s
2	349465	1m16.8s	4m16.7s
3	359659	3m01.5s	8m51.7s

As estimated by the results in Table 1, previewing by progressive refinement is approximately three times faster than previewing by ray casting without anti-aliasing. It should be added that these numbers do not entirely reflect the reality of the situation because, as demonstrated in the example of Figure 5, progressive refinement previewing already gives an accurate rendering of the surface at early stages of refinement. From a perceptual point of view, therefore, the difference between the two previewing techniques is greater than what is shown in Table 1.

Figure 6 shows two snapshots of a progressive refinement rendering where a region of interest is active. The surface being rendered is the same two layer Perlin noise surface that was shown in Figure 4. The rectangular ROI is defined on the lower right corner of the image. The portion of the surface that projects inside the ROI is given priority during progressive refinement.

Figure 7 shows the final rendering result obtained with our progressive refinement renderer for a procedural planet modelled as an implicit surface and built from a combination of different types of procedural noise functions that include Perlin noise, sparse convolution noise and cellular texture noise (Perlin, 2002; Lewis, 1989; Worley, 1996). The landscape is obtained by modulating the surface of a sphere at a very small scale.




Figure 5: From left to right, top to bottom, snapshots taken during the progressive refinement rendering of a procedural noise function. The snapshots were taken after 5000, 10000, 28000 and 350759 iterations, respectively. The wall clock times at each snapshot are 1.02s, 1.98s, 4.18s and 27.80s, respectively.




Figure 6: Progressive refinement rendering with an active region of interest shown as a red frame. Once rendering is complete inside the region, refinement continues on the rest of the image.

5 CONCLUSIONS

The rendering method, here presented, offers the possibility of visualising implicit surfaces with progressive refinement. The main features of a surface become visible early in the rendering process, which makes this method ideal as a previewing tool during the editing stages of an implicit surface modeler. In comparison, a meshing method would generate expensive high resolution preview meshes for the more complex surfaces while a ray caster would be slower and without the progressive refinement feature. Our rendering method, however, does not implement anti-aliasing and cannot compete with an anti-aliased ray caster as a production tool. Production quality renderings of some of the surfaces shown in this paper are typically done overnight, a fact which further justifies the need for a previewing tool.

It would have been straightforward to incorporate anti-aliasing into our rendering method by allowing cells to be subdivided down to sub-pixel size and then applying a low-pass filter to reconstruct the pixel samples. There is, however, one issue that prevents the use of our method for high quality renderings and which makes such implementation effort not worthwhile. As explained in Section 3.1, the computation of range estimates with affine arithmetic is always conservative. This conservativeness implies that some cells a small distance away from the surface may be incorrectly flagged as intersecting with it. As a consequence, some portions of the surface may appear dilated after rendering. The offset error at some point on the surface is in the same order as the size of a pixel times the distance to the point. This artifact can be tolerated during previewing but is not acceptable for production quality renderings.

We intend in the future to apply our progressive refinement previewing strategy not only to procedural fractal planets in implicit form but also to implicit surfaces that interpolate scattered data points.

REFERENCES

- Bergman, L., Fuchs, H., Grant, E., and Spach, S. (1986). Image rendering by adaptive refinement. In Evans, D. C. and Athay, R. J., editors, *Computer Graphics (SIGGRAPH '86 Proceedings)*, volume 20, pages 29–37. ACM Press.
- Bloomenthal, J. (1988). Polygonisation of implicit surfaces. *Computer Aided Geometric Design*, 5(4):341–355.
- Cohen, M. F., Chen, S. E., Wallace, J. R., and Greenberg, D. P. (1988). A progressive refinement approach to fast radiosity image generation. In Dill, J., editor, *Computer Graphics (SIGGRAPH '88 Proceedings)*, volume 22, pages 75–84. ACM Press.
- Comba, J. L. D. and Stolfi, J. (1993). Affine arithmetic and its applications to computer graphics. In *Proc. VI Brazilian Symposium on Computer Graphics and Image Processing (SIBGRAP '93)*, pages 9–18.
- Cook, R. L. (1989). Stochastic sampling and distributed ray tracing. In Glassner, A. S., editor, *An Introduction to Ray Tracing*, chapter 5, pages 161–199. Academic Press.
- de Cusatis Jr., A., de Figueiredo, L. H., and Gattas, M. (1999). Interval methods for raycasting implicit surfaces with affine arithmetic. In *Proc. XII Brazilian Symposium on Computer Graphics and Image Processing (SIBGRAP '99)*, pages 65–71.
- de Figueiredo, L. H. and Stolfi, J. (1996). Adaptive enumeration of implicit surfaces with affine arithmetic. *Computer Graphics Forum*, 15(5):287–296.
- Duff, T. (1992). Interval arithmetic and recursive subdivision for implicit functions and constructive solid geometry. In Catmull, E. E., editor, *Computer Graphics (SIGGRAPH '92 Proceedings)*, volume 26, pages 131–138. ACM Press.
- Ebert, D. S., Musgrave, F. K., Peachey, D. R., Perlin, K., and Worley, S. P. (2003). *Texturing & Modeling: A Procedural Approach*. Morgan Kaufmann Publishers Inc., 3rd edition.
- Farrugia, J. P. and Peroche, B. (2004). A progressive rendering algorithm using an adaptive perceptually based image metric. *Computer Graphics Forum*, 23(3):605–614.
- Guo, B. (1998). Progressive radiance evaluation using directional coherence maps. In Beach, R. J., editor, *Computer Graphics (SIGGRAPH '98 Proceedings)*, volume 22, pages 255–266. ACM Press.
- Hart, J. C., Jarosz, W., and Fleury, T. (2002). Using particles to sample and control more complex implicit surfaces. In *Proceedings Shape Modeling International*, pages 129–136.
- Laur, D. and Hanrahan, P. (1991). Hierarchical splatting: A progressive refinement algorithm for volume rendering. In Sederberg, T. W., editor, *Computer Graphics (SIGGRAPH '91 Proceedings)*, volume 25, pages 285–288. ACM Press.
- Lewis, J.-P. (1989). Algorithms for solid noise synthesis. In Lane, J., editor, *Computer Graphics (SIGGRAPH '89 Proceedings)*, volume 23, pages 263–270. ACM Press.
- Lippert, L. and Gross, M. H. (1995). Fast wavelet based volume rendering by accumulation of transparent texture maps. *Computer Graphics Forum*, 14(3):431–444.
- Lorensen, W. E. and Cline, H. E. (1987). Marching cubes: A high resolution 3D surface construction algorithm. In Stone, M. C., editor, *Computer Graphics (SIGGRAPH '87 Proceedings)*, volume 21, pages 163–169. ACM Press.
- Maillot, J.-L., Carraro, L., and Peroche, B. (1992). Progressive ray tracing. In Chalmers, A., Paddon, D., and Sillion, F., editors, *Third Eurographics Workshop on*




Figure 7: The visualisation of a procedural landscape that corresponds to a small section of an entire fractal planet. The image shows the final rendering result obtained with progressive refinement.

- Rendering*, pages 9–19, Bristol. Eurographics, Consolidation Express Publishing.
- Messine, F. (2002). Extensions to affine arithmetic: Application to unconstrained global optimization. *Journal of Universal Computer Science*, 8(11):992–1015.
- Mitchell, D. P. (1990). Robust ray intersection with interval arithmetic. In *Proceedings of Graphics Interface '90*, pages 68–74. Canadian Information Processing Society.
- Moore, R. (1966). *Interval Arithmetic*. Prentice-Hall.
- Painter, J. and Sloan, K. (1989). Antialiased ray tracing by adaptive progressive refinement. In Lane, J., editor, *Computer Graphics (SIGGRAPH '89 Proceedings)*, volume 23, pages 281–288. ACM Press.
- Perlin, K. (2002). Improving noise. *ACM Transactions on Graphics (SIGGRAPH '02 Proceedings)*, 21(3):681–682.
- Roth, S. D. (1982). Ray casting for modeling solids. *Computer Graphics and Image Processing*, 18(2):109–144.
- Stolfi, J. and de Figueiredo, L. H. (1997). Self-validated numerical methods and applications. Course notes for the 21st Brazilian Mathematics Colloquium.
- Velho, L. (1996). Simple and efficient polygonization of implicit surfaces. *Journal of Graphics Tools*, 1(2):5–24.
- Whitted, T. (1980). An improved illumination model for shaded display. *Communications of the ACM*, 23(6):343–349.
- Witkin, A. P. and Heckbert, P. S. (1994). Using particles to sample and control implicit surfaces. In Glassner, A., editor, *Computer Graphics (SIGGRAPH '94 Proceedings)*, volume 28, pages 269–278. ACM Press.
- Woodwark, J. R. and Quinlan, K. M. (1982). Reducing the effect of complexity on volume model evaluation. *Computer Aided Design*, 14(2):89–95.
- Worley, S. P. (1996). A cellular texture basis function. In Rushmeier, H., editor, *Computer Graphics (SIGGRAPH '96 Proceedings)*, volume 30, pages 291–294. ACM Press.