

## Introduction

The *Graphicsland* project group [Wyvill, B. 85a] at the University of Calgary has developed an organised collection of software tools for producing animations from models in three dimensions. The system allows the combination of several different kinds of modelling primitive [Wyvill et al 85b]. Thus polygon based models can be mixed freely with fractals [Mandelbrot 83, Fournier 82] and particles [Reeve 83] in a scene. Motions and camera paths can be described, and animations generated. Note that we do not include the use of a two dimensional 'paint' system. Our objective is always to construct views of a full three dimensional model.

An important class of objects in the everyday world is *soft*. That is, the shape of the object varies constantly because of the forces imposed on it by its surroundings. A bouncing ball is a simple example: as it strikes the ground, it flattens. The smoothly covered joints of animals change shape with seamless continuity, and liquids mould themselves to their surroundings and even break into separate droplets. Even apparently rigid objects deform in some circumstances. Trees, for example, bend in the wind.

To date, there seems to have been few attempts to model *soft* objects for computer graphics. Possibly, this is because *soft* objects are less important in engineering. But it is also true that much effort in computer graphics has been directed to producing still pictures and you cannot tell that an object is *soft* until it moves. Clouds [Gardner 85] and particles [Reeve 83] come close, but there is nothing in either of these papers which deals with the interaction of particles with surrounding objects.

We have been experimenting with a general model for *soft* objects which represents an object or collection of objects by a scalar field. That is a mathematical function defined over a volume of space. The object is considered to occupy the space over which the function has a value greater than some threshold so the surface of the object is an iso-surface of the field function. That is a surface of constant function value within the space considered.

This is not a new idea. The technique is described by James Blinn, and used to create models of molecules [Blinn 1982]. He also suggests other applications and describes a direct rendering technique using an elegant set of sorted lists. Ken Perlin has used a modification of Blinn's method to represent 'stochastic' shapes [Perlin 1985]. We, however, are interested in simpler shapes which we can *move convincingly*.

By suitable choice of field function, we can represent a wide variety of shapes conveniently and, in principle, any shape somehow. For this paper we concentrate on simple functions based on proximity to given data points. We achieve animation by specifying the motions of these key points, without otherwise altering our function.

Strictly speaking, the word 'field' refers to a particular set of values distributed across space at some time. Thus the value of 'field' at some point  $\langle x,y,z \rangle$  is found by evaluating the function  $f(x,y,z)$ . In what follows we use the words 'function' and 'field' more or less interchangeably.

In this paper we describe our choice of function and the data structures which enable us to construct the surface quickly. In the near future we will also report on the techniques for moving the key points for some simple models.

## Space Function or *Field*

We want to construct a function which will enable us to represent arbitrary shapes when we plot the iso-surfaces. The function is therefore going to depend on a set of given key points. We assume that the key points are independent. That is, we treat them like particles in a cloud. They are all alike and the field value at a point in space depends only on the proximity of key points. (These assumptions are arbitrary. One can easily imagine a function of key points which includes some knowledge about relations between points. The points could be ordered, for example, and a function value calculated by interpolation between adjacent point values.)

For our models we want the field to be continuous. For some models, we may use hundreds, thousands or even tens of thousands of key points. So it is important that we do not have to inspect every key point whenever we wish to calculate a field value. Therefore we use a function which is not influenced by any point beyond a certain distance away. This distance is known as the radius of influence,  $R$ . Again, arbitrarily, we make the field value due to several nearby points equal to the sum of the values due to each point. This is the very simplest way to combine the effect of many points and it seems to work well enough.

By definition, the contribution to the field made by any key point beyond its radius of influence,  $R$ , is zero. The contribution at the position of the point itself will be some maximum value (we use 1.0) and we would like to arrange that the field drops smoothly to zero at  $R$ . If we express this contribution to the field as a function  $C$  of  $r$ , the distance from the key point, these requirements can be expressed in terms of the values of the function at the point,  $r=0.0$ , and at the radius of influence  $r=R$ .

$$\begin{array}{ll} C(0.0) = 1.0 & C'(0.0) = 0.0 \\ C(R) = 0.0 & C'(R) = 0.0 \end{array}$$

where  $C'$  is the derivative of  $C$  with respect to  $r$ . These conditions are sufficient to define a unique cubic function for  $C$ :  $C(r) = \frac{2r^3}{R^3} - 3\frac{r^2}{R^2} + 1$  The field at any point  $\langle x,y,z \rangle$  is then the sum of  $C(r)$  calculated for each key point within  $R$ . This field function turns out to be quite satisfactory but a little slow because calculation of  $r$  requires a square root. We eliminate the need for this by using a cubic in  $r$  squared and adjusting the coefficients to make it approximate the original function  $C$ . A cubic in  $r$  squared is guaranteed to have  $C'(0.0)=0.0$ , so we can add another condition:  $C(r)=f$ , for some chosen values  $r,f$ . We use  $r=R/2$  and  $f=0.5$  and this is close enough to the original function to make no difference in practice. Our field function is shown graphically in Figure 1.

Blinn used an exponential function for his field based on the known field of electron density around an atom [Blinn 1982]. Our 'cubic' function is similarly shaped and has the desirable property of dropping to zero at the radius of influence,  $R$ . It is also very cheap to calculate, needing only three additions and three multiplications.

## Defining the Iso-surface

Having established a definition of the field, we must choose a field value for the iso-surface. Clearly the field due to a single key point will be symmetrical about the point and any iso-surface in that field will be a sphere. Suppose we choose a function value, *magic*, and plot the iso-surface connecting all points whose field value equals *magic*. Now consider the field due to two key points in the same place. It is still symmetrical and an iso-surface for value *magic* in this field will be a sphere of larger radius than the iso-surface in the field due to one point. We have chosen *magic* so that this larger sphere has exactly twice the volume of the other.

This choice is intended to suit the modelling of liquids, to provide a reasonable effect when two droplets merge. Other choices are possible as are other functions for the field. Finding functions appropriate to particular applications is a research project in its own right. For our purposes, the above is used throughout.

## Producing the Surface

The surface defined in this way, by a collection of data points, is very general. It is not even necessarily connected and in order to make a picture, we first convert it to a more tractable form. We have chosen to use a simple polygon mesh for this purpose.

We construct the mesh in two distinct stages. Imagine that the part of space occupied by the surface is divided by a three dimensional grid into small cubes. First we find all the cubes which are intersected by the surface and then we construct the polygons in each cube.

To find the cubes intersected by the surface without scanning the whole of a large three dimensional grid, we take advantage of the knowledge that all our key points are enclosed by some part of the surface. For each key point, starting at the nearest grid point, we calculate the field at a succession of adjacent grid points along one axis until we encounter a point whose field value is less than *magic*. This point and the previous one form the endpoints of one edge of a cube which is intersected by the surface. This process gives us a set of 'seed' cubes such that every disconnected component of the surface intersects at least one seed cube.

Each seed cube shares faces with six neighbours. Starting at the seed cubes, we examine each cube's neighbours to see whether or not it is intersected by the surface. If a neighbour is intersected then we look at its neighbours and so on until all of the cubes intersected by the surface have been found. This completes the first stage.

In the second stage, we have only to deal with cubes which are intersected by the surface. For each cube we have eight values which are the field values at its vertices. From these we construct a set of polygons which are part of the iso-surface. The previous stage has sorted out all the cubes intersected by the surface, so when we put these polygons together we have a representation of the surface.

To complete the description we must explain the data structure used in the first stage and the logic used in the second.

## Data Structure

There are two distinct problems, each handled by a structure using a hash table.

The first problem is to be able to calculate the field value at any point efficiently. This means that we must avoid inspecting key points which are too far away to affect the result. We do this by dividing our space into cubes, and sorting the key points according to which cubes they occupy. We only need to look at key points which occupy cubes which are close to the point where we are finding the field. The data structure for this is shown in Figure 2.

The key points are represented by a linked list of triples of floating point values,  $\langle x, y, z \rangle$ . These are grouped spatially into cubes of side  $R$ , where  $R$  is the radius of influence (defined above). All the triples in such a list belong to one group.

The group is represented by a pointer to such a list and a triple  $\langle l, m, n \rangle$ , where  $\langle l, m, n \rangle$  represents the 'low' vertex of a cube of side  $R$ , containing one or more key points. The 'low' vertex is the bottom-south-west corner, or, more formally, the vertex of lowest  $\langle x, y, z \rangle$ . The groups are also elements of a linked list.

Key points are accessed via a hash table. A hash address is calculated from the  $\langle l, m, n \rangle$  triple. This is used as an index in the table. The table entry contains a pointer to a linked list of groups. To access all the key points in a given group, we first find the list of groups from the hash table. Then search down the list of groups until we find one which matches our starting values of  $\langle l, m, n \rangle$ . This group contains the pointer to the appropriate list of key points.

Any cube in space has twenty six neighbouring cubes. To find the field value at  $\langle x, y, z \rangle$ , we must examine every key point within a distance  $R$  of  $\langle x, y, z \rangle$ . These points must lie within the same cube as  $\langle x, y, z \rangle$  or in one of its neighbours. Using this hash table and linked structure we can quickly check the twenty seven possible groups and thus calculate the field without inspecting every key point.

In doing this, we are inspecting all the key points within a volume of space:  $V=27R^3$  but all the points we are interested in lie in a sphere of radius  $R$  and volume:  $W=4\frac{\pi}{3}R^3 \approx 4.12V$

By making our groups represent smaller cubes, and searching a larger collection of 'neighbours' we can find the field by inspecting fewer points. But when the number of points in each group is small, this is wasteful because of the time taken to search more neighbouring cubes. The maximum possible saving is a factor of four and the extra work for searching more cubes goes up rapidly. In our models, the key points are not crowded within the radius of influence and the field at any point usually depends on less than thirty key points. Our simple strategy is thus justified.

Our second problem is to avoid recalculating field values as we find the cubes which intersect the surface. This is done with a second hash table, Figure 3. We are only interested in calculating the field at grid points which are the vertices of our little cubes. It is convenient, therefore, to represent these points by integers  $i, j, k$  where  $i*d, j*d, k*d$  are the actual co-ordinates of the point and  $d$  is the grid spacing. Each vertex is represented by a quintuple  $\langle i, k, j, f, done \rangle$  and these are linked in a list for access through the hash table. 'f' is the field value for the vertex and the meaning of 'done' is explained below.

To access the quintuple representing a given vertex, a hash code is calculated from its  $i,j,k$  values. This is used as an index in the table and the table entry contains a pointer to a linked list of vertices. This list must be searched to find the particular vertex. When a vertex is referred to for the first time, the search fails. In this case the field value,  $f$ , is calculated and a new vertex is linked into the list. Each vertex is shared by as many as eight cubes but its field value need be calculated only once. Subsequent references will find it stored in the structure.

To trace all the cubes intersected by the iso-surface, it is necessary to mark those cubes which have already been processed. For this purpose, each vertex is also considered to represent the cube of which it is the 'low' vertex. The flag 'done' within the vertex is made TRUE to indicate that this cube has been dealt with.

The algorithm for finding all cubes on the surface from a seed cube can now be described. Observe the numbering of the cube vertices in Figure 4. The low vertex of the cube is numbered 0 and this is the vertex  $\langle i,j,k \rangle$ . Consider the four vertices 0,4,6,2. These are shared by the cube  $\langle i-1,j,k \rangle$ . If the field values of these vertices are all greater than the iso-surface value or if they are all less, then the surface does not pass through the face 0,4,6,2. If some of the values are greater and some less than the iso-surface value, the surface does pass through the face and the cube  $\langle i-1,j,k \rangle$  must be processed if this has not already been done. To process a cube, first its 'done' flag is made TRUE and then a pointer to the low vertex of the cube is entered in a queue. More precisely:

```

begin Set seed cube's done flag to true.
  Add seed cube to the queue.
  while queue is not empty do
    begin Remove one cube from the queue.
      for each face of cube do
        begin if surface intersects face then
          begin select neighbour cube for that face.
            if neighbour's done flag is not true then
              begin Set neighbour's done flag to true.
                Add neighbour to queue.
              end
            end
            Pass vertices and values for cube to second stage.
          end
        end
      end
    end
  end
end

```

## The Hash Functions

Most of the space is empty. That is to say that the cubes which intersect our *soft* object's surface represent a small fraction of the total. Similarly, the group cubes of side  $R$  which contain key points are few compared with the number of such cubes in the region of space we are dealing with. If this were not so, we would find no advantage in using a hash table. We would simply store points in an array. So our hash function must have an even distribution for points which are geometrically close together. In both of our tables, the hash function maps a triplet to an address. For a table of size  $l*m*n$  we map  $\langle i,j,k \rangle$  to:

$$(\text{rem}(i,l)*m+\text{rem}(j,m))*n+\text{rem}(k,n)$$

where  $\text{rem}(x,y)$  is the remainder when  $x$  is divided by  $y$ . We use  $l=m=n=16$  and the function is calculated quickly by logical operations. A more sophisticated function might be appropriate for some applications but we have found no reason to change it.

## The Second Stage: Generating The Polygons

In this part of the process, we are given the field values at the vertices of a cube in the mesh and we must construct polygons which represent the part of the iso-surface which intersects that cube.

Firstly we find points which approximate the intersection of the iso-surface with the edges of the cube. Then we connect these to make the polygons. To avoid ambiguity we refer to these points as 'intersections'. The word 'vertex' is reserved for vertices of the cube. A vertex whose function value is greater than *magic*, we call hot, and a vertex whose function value is less than or equal to *magic* we call cold. Vertices whose field value is exactly *magic*, present minor complications. (We have to take special action to avoid generating polygons of zero size.)

Suppose two adjacent vertices  $p,q$  have field values  $f_p, f_q$ . If  $q$  is hot and  $p$  is cold, then  $\frac{\text{magic}-f_p}{f_q-f_p}$  is taken as the distance from  $p$  to the intersection of the iso-surface with  $pq$ . Although this 'linear interpolation' is not the true intersection, it is much cheaper to calculate and, provided the cubes are small enough that the polygonal approximation to the surface is reasonable, it is good enough. Note that this calculation is consistent across adjacent cubes which share edges.

The process by which we connect these points is fairly complex, so we start with an example. Figure 5(A) shows a cube with only one hot vertex. Clearly the intersecting iso-surface is approximated by a single triangle. In Figure 5(B), there are two adjacent hot vertices and we have a quadrilateral. In Figure 5(C), there are two non-adjacent hot vertices and the polygon is a hexagon which encloses both of them. Why do we choose this hexagon rather than two triangles, Figure 5(D)? In effect, we are electing to link the two hot vertices and separate the two cold ones. Whether to link the hot or cold vertices in this case is decided by examining the values of the field. The value at the centre of the face is approximately the mean of the four vertex values. If this value is greater than *magic* we link the hot vertices.

Each intersection is uniquely associated with one edge of the cube. So we can label an intersection by the pair of vertices  $\langle p,q \rangle$  between which it lies. The possible configurations of one face of the cube are shown in Figure 6. When the number of hot vertices in this face is four or zero, no polygon edges are created: cases A,B. When the number of hot vertices is one or three, a single edge is created: cases C,D. For two hot vertices we have case E: create one edge; case F: create two edges linking hot vertices; and case G: create two edges linking cold vertices. These created polygon edges are represented by ordered pairs of intersections. That is pairs of vertex pairs. In the diagram, the order is indicated by the arrows.

These polygon edges are stored in an array, indexed by the first intersection. The second intersection is always the same as the first intersection of some other edge, so we can form a polygon by tracing the natural successors of each edge until we return to the edge we started at.

The algorithm follows:

```

for each edge of cube, <p,q> do
  if p is hot and q is hot or p is cold and q is hot
  then create intersection <p,q>;
for each face of cube do create edges according to Figure 6;
while edges remain do
  begin start:=any edge; polygon:={start};
  remove start from edge array;
  next:=successor of start;
  while next <> start do
    begin polygon := polygon + {next};
    remove next from edge array;
  end;
  output polygon
end

```

The polygons are not, in general, planar. So we divide them into triangles by connecting each intersection to a centroid. This method of triangulation doesn't work for polygons in general, but it seems to be alright for polygons generated by this algorithm.

## Animating the Objects

The shape of a *soft* object is entirely determined by the positions of the key points. So we describe its motion and shape changes by moving key points only. A general principle we have adopted for natural animation is to use a physical simulation rather than interpolate between key frames. Traditional fairing techniques can be viewed as an attempt to do this without performing the calculation needed for simulation. In this paper we describe only a few examples.

## Examples Of Application

Figure 7 shows two droplets merging in stages. This is quite pleasing in animation. The droplets are modelled by single key points with  $R=2.0$ . The distance,  $r$ , between the two key points is shown for each stage.

Figures 8 to 10 are frames from an animation of a *soft* object sliding down steps. The object is roughly shaped as the letters of *Soft* and these distort smoothly as they slide down.

Figures 11 and 12 show the letters of *Soft* rising from a trough of bubbling *soft* material. The background in Figure 11 features fractal mountains and rolling hills. These are part of the solid scene and show how the *soft* objects have been incorporated into the *Graphicsland* system.

## Discussion

Our modelling technique has proven to be a convenient tool in computer animation. So far, we have concentrated on one very simple method of making our field functions and there are many other possibilities which are worthy of investigation. The field can be generated by arbitrary functions permitting the modelling of mathematical surfaces. Our field function can be modified by extra terms which are not related to the key points. This can be used to make objects deform in special ways. For example, if we add a surface held to a negative field value, objects can be made to vanish bit by bit as they approach the surface, and to reappear on the other side. If we make a 'hole' in the surface, an object approaching the hole, appears to squeeze through it.

We have represented our surfaces by polygon patches. This is because the *Graphicsland* system already offered us versatile display tools for this kind of model. It would clearly be a good idea to use bicubic or conic patches instead and we are planning some experiments to do this. We have not followed Blinn's technique [Blinn 1982] of rendering the surface directly, because we wanted to produce objects which we could easily incorporate into *Graphicsland*. Also, our models can contain a very large number of key points and most of them are not near to any surface. This is particularly true of models of liquids. The polygon generating algorithm is particularly efficient in this case because the 'hidden' key points do not have to be known to the rendering algorithm.

We are, however, experimenting with a ray-tracer which renders the field directly. In this case, we use our *surface following* algorithm to construct a boundary of cubes which contains the iso-surface.

Even when we have a superior rendering technique for these surfaces, an efficient generator of polygon patches will be useful. *Graphicsland* does not yet offer interactive graphical development of animations, because it takes too long to render the scenes. We are introducing a new feature to produce very fast rendering of scenes at greatly reduced quality. For this purpose we expect to continue to need a polygonal representation.

Our *soft* objects are reasonably quickly generated. The letters and contents of the trough in Figure 11 took about eight minutes per frame on a VAX11/780 computer to generate the polygons. This is much less than the rendering time for 512 by 512 pixels using a z-buffer algorithm.

## Acknowledgements

The *JADE* project at the University of Calgary has been particularly supportive. This work and *JADE* is supported by the Natural Science and Engineering Research Council of Canada.

## References

- Blinn, J (1982) "A Generalization of Algebraic Surface Drawing" *ACM Transactions on Graphics*, 1, 235.



- Fournier, A, Fussel, D, and Carpenter, L (June 1982) "Computer Rendering of Stochastic Models" *Commun. ACM*, 25, 6, 371-384.
- Gardner, G. (July 1985) "Visual Simulation of Clouds" *SIGGRAPH 85, Computer Graphics*, 19 (3) 297-303.
- Mandelbrot, Benoit. (1983) "The Fractal Geometry of Nature." *W.H. Freeman and Company. (First Edition 1977).*
- Perlin, K. (July 1985) "An Image Synthesizer" *SIGGRAPH 85, Computer Graphics*, 19 (3) 287-296.
- Reeves, William. (Apr 1983) "Particle Systems - A Technique for Modeling a Class of Fuzzy Objects" *ACM Transactions on Graphics*, 2, 91-108.
- Wyvill, B.L.M., McPheeters, C., and Garbutt, R. (July 1985a) "A Practical 3D Computer Animation System" *The BKSTS Journal (British Kinematograph Sound and Television Society)*, 67 (6) 328-332.
- Wyvill, B.L.M., McPheeters, C., and Novacek, M. (June 1985b) "Specifying Stochastic Objects in a Hierarchical Graphics System" *Proc. Graphics Interface 85, Montreal.*




Figure 1  
Soft Field Vs. Radius r

Hash Table

Groups

Key Points




Figure 2

Data Structure for efficient access to key points.

Hash Table

Vertex Quintuples

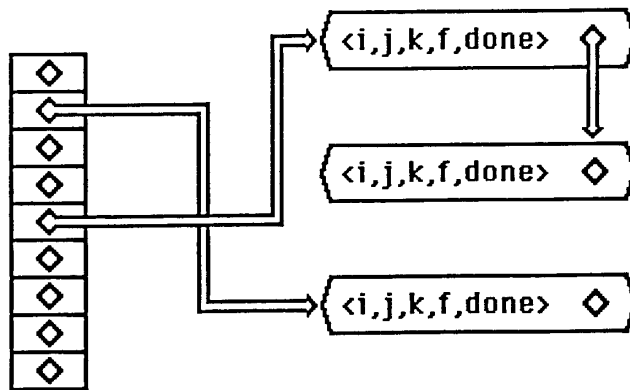


Figure 3

Data structure for cube vertices.




Figure 4

Neighbouring cubes.







Figure 5

Cubes with one and two hot vertices.




Figure 6

The seven different cases for connecting intersections.