

SHADING THREE-DIMENSIONAL IMAGES

This invention relates to a system for shading three-dimensional images and in particular to a system which can do this substantially in real time.

The best known prior art systems for generating fully occulted real time 3D images is the Z-buffer or depth buffer image precision algorithm. This may be implemented in either software or hardware and is used in systems produced by companies such as Silicon Graphics, Evans and Sutherland, and Hewlett Packard.

The Z-buffer algorithm requires a frame buffer in which colour values are stored for each pixel in an image. In addition to this it requires a Z-buffer with an entry for each pixel. In this a Z value (depth value) is stored for each pixel. To generate a 3D representation polygons, such as triangles, are rendered into the frame buffer in arbitrary order. During scan conversion of the polygon if a point on a polygon is no further from the viewer than the point already in the buffer for that pixel then the new point's colour and Z value replace the old values. No presorting is necessary and no object to object comparisons are required.

Z-buffer type systems require a high level of integration and performance and are very expensive. It is necessary to use a high performance frame buffer and Z-buffer if real time performance is to be achieved.

Preferred embodiments of the present invention enable a highly integrated, low cost, high performance real time three-dimensional graphics system to be produced.

Preferably the invention is embodied in a scalable regular pipeline structure which is highly suitable for cost effective VLSI (very large scale integration) implementation. Using such a scalable approach means that a 3D graphics rendering system may be provided on a single chip.

Refinements to the system enable features

such as shadows, anti-aliasing, soft shadows, transparency, depth of field effects, motion blur, curved surfaces, and direct implementation of constructive solid geometry (CSG) to be supported. Some of these features, e.g. real time shadow generation, are extremely difficult to implement using Z-buffer techniques and therefore very few high performance systems actually implement them. The approach taken by the present invention enables features such as shadows to be implemented in relatively straightforward manner.

The present invention is based on the use of a ray casting technique for rendering of three-dimensional images rather than conventional polygon based rendering techniques.

In a preferred embodiment of the invention objects are each represented by a set of surfaces which are stored as sets of data. An image plane is deemed to lie between a viewer and the scene to be viewed and this is comprised of a plurality of elementary areas (pixels). A ray is assumed to pass from the viewpoint through an elementary area of the screen into the scene to be viewed and will intersect various surfaces which represent an object in the scene. By analysis of these intersections and their distances from the viewer we can determine whether any surface is visible and thus whether that pixel should be rendered with a particular shade.

Using this approach requires a continuous stream of similar calculations to be performed. The set of surfaces for each object in the scene have to be tested at each pixel and the shade assigned to the closest visible surface is given to that pixel. Such a technique can be implemented in a pipeline fashion. This enables it to be highly integrated and trade-offs can be introduced between performance/complexity and cost.

The invention is defined in its various aspects with more precision in the appended claims to which reference should now be made.

An embodiment of the invention will now be described in detail by way of example with reference to

the accompanying drawings in which:

Figure 1 shows a schematic block diagram of a system embodying the invention;

Figure 2 a), b), and c) shows shapes which may be manipulated and rendered by an embodiment of the invention;

Figure 3 shows a cross-section through two ray/surface intersections of an object illustrating an embodiment of the invention;

Figure 4 shows two cells of a pipeline processor embodying the invention;

Figure 5 shows a block diagram of a processor embodying the invention;

Figure 6 shows a system using a processor embodying the invention; and

Figure 7 shows a block diagram of the architecture of the processor of Figure 6.

Figure 1 shows a 3D image synthesis pipeline 2 in which manipulation of image surfaces is performed according to a predetermined set of rules to allocate shades to individual picture elements (pixels). The three-dimensional image data which is acted upon by the pipeline 2 is stored in a 3D database 4. A geometry processor 6 is provided between the database and the pipeline so that data from the database can be manipulated prior to rendering by the pipeline. Manipulation can consist of functions such as object rotation, translation, viewpoint changes and shading calculations. Some form of user input can be provided to this so that real time interactive manipulation of the 3D database is possible. The image provided by the geometry processor to the pipeline is surface data which is converted by the pipeline to pixel data for display.

The pipeline processor of Figure 1 is based upon a ray tracing algorithm rather than a polygon based rendering algorithm. In the ray tracing algorithm each three-dimensional shape to be represented is modelled as a collection of objects as shown in Figure 2A. These may intersect one another as shown in Figure 2B in dependence

on what is required to produce an approximation of the original shape. Any concavities in the original object are broken down into a number of blocks as illustrated in Figure 2C.

Each object is defined solely as a number of surfaces. In the examples shown in Figure 2 these surfaces are all planar but the system is versatile enough to handle curved surfaces. The edges and the vertices of each object are implicitly defined from the intersections of the surfaces and thus there is no requirement to store edge or vertex data.

Before rendering of a scene can commence each surface in that scene is classified as either a forward or reverse surface. This is done by setting a flag which is stored with data defining each surface. The forward surface represents the condition where a line normal to the surface points towards the observer. The reverse surface is the opposite.

The ray tracing technique operates by taking each object in turn, and, by imagining a ray from the observer's eye passing through a chosen point (a pixel) in the image plane, intersects all the surfaces defined in that object. Once this has been done the forward surface intersection furthestest from the eye and the reverse surface intersection closest to the eye are determined. This is illustrated in Figure 3.

In Figure 3 ray 1 and ray 2 both start from an observation point 8. The object is defined by four surfaces in the plane of the cross-section of Figure 3. The forward surface intersection of Ray 1 furthestest from the observation point 8 is F1 and the reverse surface intersection closest to the observation point is R1.

This condition where the closest reverse surface intersection is nearer to the observation point than the furthestest forward intersection indicates that ray 1 does not intersect the object. This particular object is therefore not visible at the pixel in the image plane through which ray 1 passes. Ray 2, on the other hand, has its nearest reverse intersection at R4 which is

further from the observation point 8 than the furthest forward surface intersection at F3. This indicates that the object is visible in the image plane at the pixel through which ray 2 passes. This pixel is therefore allocated a shade which has been allocated to that surface. This shade is usually stored with data defining the surface and may be a grey scale value or an R, G, B value for a full colour system. This technique for allocating a shade to a particular pixel for a single object is called the inter-rule set.

The process is repeated for each object in the scene and the visible surface of the object nearest to the eye, if any, which intersects a ray passing through a particular pixel is selected to derive a shade for that pixel. This is termed the intra-rule set.

The shade which is applied to a particular pixel can then be calculated from the properties and orientation of the visible surface. If there are no visible surfaces intersected by a ray then the shade assigned to the pixel through which that ray passes will be the background shade. The shade may be a grey scale value or it may be a full R, G, B, value. It will be appreciated that the entire ray tracing process has to be executed for every point in the image plane with every surface of every object.

The process of finding surface ray intersections can be reduced to a very simple operation from the data which defines each surface. This is explained below.

Each surface can be represented in vector form in the following manner:

$$n \cdot x = k$$

where n = a unit normal vector to the surface pointing towards the origin (usually the observation point), usually stored as the cosine of the angle from the origin.

k = perpendicular distance to the origin from the surface

x = a point on a surface (coordinates)

A line or ray is represented in vector form as follows:

$$P1 + u(p2-p1)$$

where P1 = base vector or distance from the origin to the observation point.

P2 = A point on the ray in the centre of a pixel on the image plane.

U = A measure of how far along the ray an intersection occurs at point x.

Therefore:

$$n(p1+u(p2-p1)) = k$$

$$\text{and } u = \frac{k-n.p1}{n.(p2-p1)}$$

The parameter u is negative behind the surface that is intersected, ranges from 0 to 1 between p1 and p2, and is greater than 1 for points in front of the image plane. Thus, the parameter u gives a measure of distance from the observer, through a point on the image plane, to a given surface. This requires a division per pixel as we pass a ray through each pixel on the image plane. However, the inverse of this parameter, 1/u, does not require any divisions but still gives a measure of the distance to the surface. In this case it approaches zero as the distance to the surface goes towards infinity. This inverse measure of distance is normally called depth. Therefore by comparing the depth of every surface at each pixel using the following equation, we can find which surface is visible at each pixel

$$1/u = n(p2-p1)/(k-n.p1)$$

Expanding the vectors in this equation gives the following result:

$$1/u = (n1.(xs-xe) + n2.(ys-ye) + n3.(zs-ze)) / (k - (n1.xe + n2.ye + n3.ze))$$

where x_s y_s z_s are image plane intersection coordinates and x_e y_e z_e are coordinates of the eye point.

If the centre of the projection, the eye point, is also at the centre of the image plane the above equation reduces to the following:

$$1/u = (n1.x_s + n2.y_s + n3.(z_s - z_e)) / (k - n3.z_e)$$

This can be rationalised into a more general form as shown below:

$$1/u = A.x_s + B.y_s + C$$

A, B, C, are constant for each surface in a particular scene, and thus the above equation can be used to derive the distance to each surface on each ray.

x_s and y_s are coordinates of the pixel in the image plane from which the distance to each surface is to be calculated using the equation.

The image is normally built up by scanning the image plane in a raster type fashion thus for any particular surface, after the starting depth has been calculated, the depth at each pixel can be calculated by incrementing by A as we traverse in a horizontal direction in the image plane, and by B as we traverse in the vertical direction. As was stated earlier the depth of every surface has to be calculated for every pixel in the image plane. This is a considerable number of calculations. However, as the depth calculation between horizontally adjacent pixels only requires a single addition and multiple pixels can be operated on in parallel, a very efficient parallel architecture can be devised. One possible architecture which can also process shadows and transparent objects is described below with reference to Figure 4.

The inner core of a processor which embodies the invention consists of a series of identical cells which are pipelined together. A typical number would be between 32 and 64 cells. Two such cells are shown in

Figure 5. Each cell works on a single pixel and adjacent cells work on adjacent pixels, i.e.

CELL N processes Pixel(x,y)
CELL N+1 processes Pixel(x+1,y)
CELL N+2 processes Pixel(x+2,y).

This is achieved by the top portion of a cell as shown in Figure 4. For a particular surface a parameter A is lodged in a register 10 at the start of each cycle and a depth parameter representing $Ax + By + C$ is lodged in a register 12 at the start of each cycle. These two values are combined in an adder 14 to give a new depth parameter $2Ax + By + C$ for the horizontally adjacent pixel, which is lodged in the register 12 in the following cell. Parameter A is also lodged in a corresponding register 10 in the next cell. The values in registers 10 and 12 in the following cell are then combined in a corresponding adder 14 whose output forms the input to register 12 in the next cell in the pipeline.

The surfaces defining the various objects are stored as a list and are processed in a sequential fashion. The surfaces are grouped in terms of the objects which they define and are so ordered that all the forward surfaces are based at the start of the group followed by the reverse surfaces. This is shown below:

object N+1 surface M+9 reverse
object N+1 surface M+8 reverse
object N+1 surface M+7 forward
object N+1 surface M+6 forward
object N surface M+5 reverse
object N surface M+4 reverse
object N surface M+3 reverse
object N surface M+2 forward
object N surface M+1 forward
object N surface M forward.

The parameters A, B and C which define each

surface are stored with an operation code which includes other surface attributes. One of these may be a surface tag used as an index for a look-up table to shade the surface. Alternatively the tag could comprise the actual R, G, B values for that surface. Additional information in the operation code is used by the processor to decide what to do with each surface. The type of information the code might contain is as follows:

- start of a new object
- forward or reverse surface
- start of surface data for the whole scene
- end of surface data for the whole scene
- shadow object
- transparent object.

This information is represented by the control signals which are input to a signal and are stored on a register 20 in a cell. From where they may be assessed as required.

At the start of a new object the depth of the first surface in that object for the pixel being processed is stored in the intra-depth register 18. Surface tag data such as surface shade are stored in the intra-tag register 22. At the same time the previous intra-depth value from the previous object is compared against the value in the inter-depth register 16 which stores the processed visible surface depth for the pixel being processed. If the intra-depth value is greater than the inter-depth for that pixel then the value in the intra-depth register is transferred to the inter-depth register 16. Thus the system compares, with the end of the processing of the surfaces of an object, the nearest visible surface in that object with the nearest visible surface detected so far in the scene for that particular pixel. Register 16 is updated if necessary.

In order to compare intra-depth and inter-depth values MUX 24 is used to switch the output of register 16 to one of the inputs of comparator 26, the

other of which receives the output of register 18.

While the surfaces of a single object is being processed MUX 24 is controlled by the control signal lodged on register 28 to connect register 12 and register 18 to the inputs of comparator 26. For forward surfaces, if the current surface depth is greater than the intra-depth value then the intra-depth register remains the same. For reverse surfaces the opposite action is taken. Thus at the start of processing of an object the intra-depth register contains the nearest surface to the eye of the previous object. After all objects have been processed the inter-tag and the inter-depth registers contain data defining the nearest surface to the eye for the pixel that the cell was working on.

A rule generator 30 is used to implement updates to inter-depth register 16 and inter-tag register 32 when a nearer visible surface is detected.

A shadow may be rendered by the system by flagging certain objects as shadow volumes. The shadow volume is a region of space which is swept out by the silhouette edge of an object illuminated from a point light source. Any objects within this volume are deemed to be in shadow.

Shadow objects are processed in the same way as solid objects. When all shadow and non-shadow objects have been processed it is straightforward to determine whether the surface assigned to a pixel is in shadow by comparing the position of that surface with the furthest forward surface and nearest reverse surface of the shadow object and checking to see if the visible surface is within that range. If it is a flag associated with that surface is set to indicate that a modification to its shade or colour needs to be made to simulate the effect of shadowing. The system can be made versatile enough to deal with illumination from more than one point light source and thus with more than one shadow shade.

In the circuitry of Figure 4 a shadow object is handled slightly differently in that after all the forward surfaces have been processed the intra-depth

register 18 is compared with the inter-depth register 16. If the intra-depth value is greater than the inter-depth value then a flag is set in the inter-tag register 32. After all the reverse surfaces have been processed the intra-depth value is again compared with the inter-depth value. If the intra-depth value is greater than the inter-depth value then the flag is cleared. One flag is provided for each light source which can cast shadows. If after processing all the shadow objects there are any shadow flags set then the colour of the surface for that pixel will be modified to simulate the shadow condition. All these actions are performed by the rule generator 30 in Figure 4. This receives all the control signals and, via competent selectorial logic selectively enables the registers for updating.

The system is able to deal with transparency and particularly with absorption transparency. This is the situation in which light rays lose energy in particular wavelengths when passing through a transparent object and thereby cause the light ray to change colour.

The absorption effect is stimulated by flagging certain objects as being transparent. These are processed after the opaque objects. Transparent objects are processed by the usual set of intra rules. If for a particular pixel a transparent object appears in front of the visible surface at that pixel then a flag is set to indicate this condition. If the flag is set the colour of the surface at that pixel is modified to simulate transparency. Further software can be provided to derive the thickness of the transparent object thus making an appropriate further modification to the colour of that pixel.

This method is used for objects which are completely transparent. However it is possible to have a single surface of an object invisible, for example a tube. Thus non-closed convex volumes may be generated. To generate this type of arrangement, after an object has been processed the reverse surfaces are passed in again. If the forward surface was visible at a pixel and has an

invisible flag set then it has its depth and tag values replaced by the reverse surface which is nearest to the viewpoint. It is then classed as a forward surface for all further operations. If the reverse surface was also invisible then the surface is left as a reverse surface and will take no further part in depth operations. Thus open ended cylinders may be generated.

In the circuitry of Figure 4 when a transparent object is processed the intra rules stay the same. However, when a transparent intra-depth is compared against the inter-depth register the inter-depth always remains the same. If transparent intra-depth is greater than the inter-depth value then a flag is set in the inter-tag. This can then be used to modify the colour at a later stage to simulate transparency.

Figure 5 shows a simple processor including a plurality of the cells of Figure 4 which are indicated at 40. A total of M cells are cascaded together and typically M will equal 32. A memory 42 stores A, B and C values representing every surface in an image to be displayed. A register 44 stores the X and Y values for the start of each block of an image to be processed.

At the start of processing A, B and C values of the first surface in the first object to be displayed are supplied at outputs from the memory. The A value is multiplied by the initial X coordinate or the first pixel in the block in a multiplier 46 and is provided as an input to an adder 48. The corresponding B value is multiplied by the Y value of the first row of pixels to be processed in a multiplier 50 the output of which forms another input to the adder 48 the C value is supplied directly to an input of adder 48.

The output of the adder represents the depth value for the first surface at the first pixel. This is supplied to the first of cells 40.

In the first of cells 40 the inter-depth and intra-depth register 16 and 18 are both initially set to 0 and so this initial depth value will be read to both registers. The A parameter is also supplied directly to

cell 40 where it is combined with the initial depth value in the adder 14.

On the next clock cycle the A parameter is fed through to register 10 in the second cell and the updated depth value supplied by adder 14 is supplied to register 12. At the same time the A,B and C parameters for the next surface in the object being processed are supplied by multipliers 46 and 50 to adder 48 and hence the depth values are supplied to the first cell which will compare the depth of the second surface with that of the first surface and, in dependence on the comparison with the current value stored in the intra-depth register 18, and whether the new and existing surfaces are forward or reverse, may update the value stored in the intra-depth register 18. If the register is updated then the intra tag register 22 which stores surface attributes will also be updated.

At the end of the list of surfaces for an object the value in the intra-depth register is compared with that in the inter-depth register to see if the latter should be updated with a closer visible surface.

The process continues for all surfaces in the object until all have passed through the array of cells 40. As soon as the last surface in an object has been supplied by the memory 42 the first surface of the following object is supplied. Thus there is no break in the chain of surfaces being supplied to the cells 40.

Once all the surfaces for all the objects have been processed the X portion of the register 44 is incremented by M, e.g. 32, the number of cells in the processor. The entire list of objects is then again sequentially processed for the new block of cells. This process continues until an entire line of pixels has been processed at which point the Y portion of register 44 will be incremented by 1 and the X portion will be reset to its initial value representing the start of the next line of data to be processed. The connections for incrementing the X and Y portions of register 44 are not shown in Figure 5.

Various extensions to the system are possible to enhance the quality of the output picture. These are discussed below.

Anti-aliasing

Anti-aliasing can be achieved very simply by using a multi-pass technique. In this method each object is processed a number of times. Each time the object is processed its position is changed by a sub-pixel amount and the pixel colour is accumulated over the passes. At the end of the passes the final colour value is divided by the number of passes. The number of passes made depends upon the quality of the final image required. For instance, if $1/2$ pixel accuracy is required four passes would be made, with the objects moved by half pixel increments each time. Thus the ray would effectively be projected through each quarter of the pixel and the results averaged.

Soft Shadows

Soft shadows are caused by the consequence of using lights which are not point sources, thus a point on a surface may see only part of a light. The transition from seeing a partially obscured light, corresponding to the penumbra, to complete obscuration, corresponding to umbra causes the soft edge of the shadow and anti-aliasing of the shadow edge is required.

This effect can be simulated by a multipass technique where the light point used to generate the shadow volume is moved between passes. The results are accumulated between passes and the final colour value is divided by the number of passes made.

Depth of field

Depth of field, where objects are in focus or blurred as a function of their distances from the camera, simulates a camera view, rather than a human eye view. This effect can also be achieved by using a multipass approach. To achieve this the A & B parameters are slightly changed between passes and the results accumulated as before.

Motion Blur

Motion blur simulates the effect of relative

motion between a camera and a scene and is a practicable method of temporal anti-aliasing. This effect is also generated by a multipass approach. The object which is to be blurred is moved between passes and the results are accumulated as in the previous cases.

Curved Surfaces

It is possible to extend the algorithm to process curved surfaces by using the technique of forward differencing. Rather than having just three parameters associated with each surface there would now be for instance second and third order derivatives which would be accumulated as well. This would allow quadratic and cubic surfaces to be represented and thus a more realistic image can be produced.

Constructive Solid Geometry

With extensions to the rule generator and extra registers to store both the forward and reverse depths it is possible to implement some of the CSG boolean primitives. This would allow for instance, two objects to be subtracted from one another. Thus, you could generate an object with a hole in it by subtracting a smaller cube from it.

Eye Point Intersections

In the situation where a surface is oriented so as to pass directly through the eye point a divide by zero condition exists, a condition which occurs at silhouette edges. Hence the surface should go from positive infinity to negative infinity across a single pixel. As this is not possible with the system described the surface has to be perturbed to avoid this condition. This perturbation can, however, cause the silhouette edge formed by this surface to appear to move, thereby reducing the realism of the image.

By attaching a flag to the surface to indicate that it passes through the eye point it is possible to remove this problem. When the flag is asserted and the surface has positive depth then the depth is set to the maximum value allowed by the system. When the flag is asserted and the surface has negative depth

then the depth is set to the minimum depth values allowed by the system. Thus a surface can go from maximum depth to minimum depth in one pixel.

Performance Enhancements

The performance of the proposed architecture can very easily be increased by increasing the number of cells. In the previous examples 32 cells have been used. However there is no reason why 256 or more cells could not be used instead. The total number of cells will be a trade off between cost and performance given the technology which is to be used. It is also very easy to cascade a number of the proposed devices together to increase performance. This could be done by pipelining or running several devices in parallel, or a combination of both.

Another method of increasing performance would be to subdivide the screen into a number of tiles. Then, for each tile those objects within the tile would be processed, thus decreasing the average number of surfaces to be processed.

A further method which could be combined with those above would be to discard objects when it has been detected that they are no longer required in the scene. All that is required is to have a flag which is set once the view is inside an object, and which is cleared when the view is no longer in the object. When the flag is cleared that object takes no further part in the processing for that raster line.

System Architecture

A system using the architecture of the invention is shown in figure 6. This comprises a general purpose microprocessor 52 which communicates with an interface such as a games port. This communicates via a 32 bit bus with a plurality of processors 54 embodying the invention. The first processor to which the bus is connected is also coupled via two busses to a parameter store and main memory in which all the surface coordinates of a scene are stored. This may be a random access memory which can be updated with data relating to

different scenes.

The main core of each processor 54 is the array of cells described with reference to Figures 4 and 5 and is shown in more detail in Figure 7.

Figure 7 comprises a memory controller 56 which interfaces with the parameter store and main memory and with the processor port of microprocessor 52. The cell array 40 multipliers 46 and 50 and the adder 48 correspond to those shown in Figure 5 with the X and Y starting values being supplied by a control register 58. The control register 58 also supplies control signals to the memory controller which in turn supplies appropriate control signals with the surface attributes to the cells 40.

A first in first out (FIFO) register 60 is coupled to the output of the cells 40 via a multiplexer 62 and is used to accumulate processed values and thus decouple the output of the processor from the cells. A pixel accumulator 64 is provided to allow the various multipass techniques described above to be implemented.

The output of Figure 7 can be connected to a high performance bus for feeding pixel data to a frame buffer, or alternatively it could be connected directly to a digital to analogue converter and supplied to a display unit. Input to the unit can be either a games type port which would interface to joysticks, ROM cartridges, and other games input devices, or it could be a host computer interface such as IBM PC bus interface.

Memory controller 57 is able to perform functions such as rotations, translations, etc on the scene being processed in response to signals supplied from the processor port.

The whole of the circuitry shown in Figure 7 can be supplied on a single chip. Alternatively only a portion of this circuitry could be provided on a single chip.



CLAIMS:

1. A method for shading a three dimensional image for display on a screen comprising the steps of representing each object in the image by a group of infinite surfaces, for each elementary area of the display projecting a ray into the 3-dimensional image, determining the location of the intersection of the ray with each surface, determining whether any intersected surface is visible at that elementary area, and shading the elementary area for display in dependence on the result of the determination.
2. A method for shading a three dimensional image for display on a screen according to claim 1 in which each surface is designated as a forward facing surface or a reverse facing surface and the relative locations of intersections with forward and reverse facing surfaces are used in determining whether a surface is visible.
3. A method for shading a three dimensional image for display on a screen according to claim 2 wherein an elementary area is visible if the furthest forward facing surface intersection is closer to the projection point of the ray than the furthest reverse facing surface intersection.
4. A method for shading a three dimensional image for display on a screen according to any preceding claim in which for each elementary area a ray is projected in turn through the groups of surfaces representing the objects in the image.
5. A method for shading a three dimensional image for display on a screen according to claim 4 in which for each elementary area a ray is projected through each surface in turn of each group of surfaces representing an object.
6. A method for shading a three dimensional image for display on a screen according to claims 4 or 5 in which the nearest visible surface in a group of surfaces representing an object is compared



with the nearest visible surface from all preceding groups of surfaces and the elementary area is shaded in dependence on the result of the comparison.

7. A method for shading a three dimensional image for display on a screen according to any preceding claim including the step of designating an object as a shadow object, determining the boundary of the shadow object on each ray passing through each elementary area and shading an elementary area with a modified shade if a surface visible in that elementary area is within the boundary of the shadow object.
8. A method for shading a three dimensional image for display on a screen according to any preceding claim including the step of designating an object as a transparent object and modifying the shade applied to an elementary area if the visible surface at the elementary area is behind the transparent object in the image.
9. Apparatus for shading a three dimensional image for display on a screen comprising means for supplying data defining a group of infinite surfaces representing each object in the image, means for generating a depth value for each surface for each elementary area of the display in dependence on the distance of that surface from an image plane, means for determining whether any surface is visible at that elementary area, and means for shading that elementary area in dependence on the result of the determination.
10. Apparatus for shading a three dimensional image for display on a screen according to claim 9 in which each surface is designated as a forward facing surface or a reverse facing surface and the relative depths of forward and reverse facing surfaces are used in determining whether a surface is visible at an elementary area.
11. Apparatus for shading a three dimensional image for display on a screen according to claim 9 or 10, including means for comparing the relative depths of visible surfaces in the image for each elementary area and shading the elementary area in dependence on the



result of the comparison.

12. Apparatus for shading a three dimensional image for display on a screen according to claim 9,10 or 11 in which the depth value generating means generates a depth value for each surface of each object in turn for each elementary area of the image.

13. Apparatus for shading a three dimensional image for display according to claim 9,10,11 or 12 provided on a single integrated circuit.

14. A method for shading a three dimensional image for display comprising the steps of supplying data defining a group of infinite surfaces representing each object in the image, for each elementary area of the display generating a depth value for each surface in dependence on the distance of that surface from an image plane, determining whether any surface is visible at that elementary area, and shading that elementary area on dependence in the result of the determination.

15. A method for shading a three dimensional image for display on a screen according to claim 14 in which each surface is designated as a forward facing surface or a reverse facing surface and the depth values of the forward and reverse facing surfaces are used in determining whether a surface is visible at that elementary area.

16. A method for shading a three dimensional image for display on a screen according to claim 15 wherein an elementary area is visible if the depth value for the furthest forward facing surface is less than the depth value for the furthest reverse facing surface.

17. A method for shading a three dimensional image for display on a screen according to any of claims 14, 15, and 16 in which for each elementary area depth values are determined in turn through the groups of surfaces representing the objects in the image.



18. A method for shading a three dimensional image for display on a screen according to claim 17 in which for each elementary area a depth value is determined for each surface in turn of each group of surfaces representing an object.

19. A method for shading a three dimensional image for display on a screen according to claim 17 or 18 in which the nearest visible surface in a group or surfaces representing an object is compared with the nearest visible surface from all preceding groups of surfaces and the elementary area is shaded in dependence on the result of the comparison.

20. A method for shading a three dimensional image for display on a screen according to any of claims 14 to 19 including the step of designating an object as a shadow object, determining the boundary of the shadow object and shading an elementary area with a modified shade if a surface visible at that elementary area is within the boundary of the shadow object.

21. A method for shading a three dimensional image for display on a screen according to any of claims 14 to 20 including the step of designating an object as a transparent object and modifying the shade applied to an elementary area if the visible surface at the elementary area is behind the transparent object in the image.

22. A method for shading a three dimensional image for display on a screen substantially as herein described.

23. Apparatus for shading a three dimensional image for display on a screen substantially as herein described with reference to the accompanying drawings.