



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 7 : G06T 15/00		A1	(11) International Publication Number: WO 00/28477
			(43) International Publication Date: 18 May 2000 (18.05.00)
<p>(21) International Application Number: PCT/GB99/03716</p> <p>(22) International Filing Date: 8 November 1999 (08.11.99)</p> <p>(30) Priority Data: 9824406.4 6 November 1998 (06.11.98) GB</p> <p>(71) Applicant (<i>for all designated States except US</i>): IMAGINATION TECHNOLOGIES LIMITED [GB/GB]; Home Park Estate, Kings Langley, Hertfordshire WD4 8LZ (GB).</p> <p>(72) Inventor; and</p> <p>(75) Inventor/Applicant (<i>for US only</i>): GIBSON, Cliff [GB/GB]; 224 Cotterells, Hemel Hempstead HP1 1JP (GB).</p> <p>(74) Agent: ROBSON, Aidan, John; Reddie & Grose, 16 Theobalds Road, London WC1X 8PL (GB).</p>			(81) Designated States: JP, US, European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).
			Published <i>With international search report.</i>
<p>(54) Title: IMAGE PROCESSING APPARATUS</p> <p>(57) Abstract</p> <p>Image processing apparatus (60) for rendering (i.e. coloring, texturing or shading) an image includes a tiling device (66) which divides the image into sub-regions or tiles. Two rendering devices (70A, 70B) are provided, and the tiles are allocated so that some are processed by one rendering device and some by the other. Polygons representing surfaces of objects to be displayed are tested against the tiles. If the surface falls into one sub-region only, the data is sent to one rendering device only. On the other hand, if the surface falls into two sub-regions being handled by the different rendering devices, then the data is sent to both rendering devices. The result is that a substantial proportion of the data need only be supplied to and processed by one rendering device, thereby speeding the operation of the apparatus. The outputs of the two rendering devices (70A, 70B) are subsequently combined by tile interleaving and image display circuitry (72).</p>			

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

- 1 -

IMAGE PROCESSING APPARATUS

Background of the Invention

This invention relates to apparatus for the processing of images.

5 The invention is particularly, though not exclusively, suitable for use in systems for the real-time rendering, texturing or shading of three-dimensional (3D) images. Real-time here means sufficiently quickly for the image to be displayed without appreciable perceptible
10 delay to the viewer.

The best known existing system for generating real-time 3D images is the Z-buffer (or depth buffer) image precision algorithm. The Z-buffer algorithm requires a frame buffer in which color values are stored for each
15 pixel (elementary picture element) in an image. In addition to this it requires a Z-buffer with an entry for each pixel. In this Z-buffer, a Z value or depth value is stored for each pixel. To generate a 3D representation, polygons are rendered into the frame buffer in arbitrary
20 order. As a subsequent polygon is entered into the frame buffer, if a point on the polygon is nearer to the viewer than the point already in the frame buffer for that pixel, then the new point's color and Z value replace the previously-stored values. If texture or shading of the
25 polygon is desired, the texturing or shading is applied to the polygon before it is rendered into the frame buffer. The system is fed with a list of vertices which define the polygons, and texture and shading operations are performed on each polygon before a depth test is executed. The
30 performance of such systems is limited by various factors, including the input data bandwidth, the speed of texturing and shading, and the local memory interface bandwidth. It has been proposed to improve the system performance by the

- 2 -

use of dual rendering devices, which are 'scan line interleaved'. That is, there are two processors which process alternate scanlines of the image raster, thus sharing the processing between them.

5 Another system for generating real-time 3D images is described in United States Patent US-A-5,729,672 assigned to VideoLogic Limited. This system uses a 'ray-casting' technique for the rendering of three-dimensional images rather than conventional polygon-based rendering
10 techniques. In this system, 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 image plane is composed of a plurality of pixels. A ray is assumed to pass from the
15 viewpoint through a pixel of the screen into the scene to be viewed and will intersect various surfaces which represent objects in the scene. By analysis of these intersections and their distances from the viewer, the system can determine whether any surface is visible. If
20 it is visible, that surface is then textured or shaded as desired. However if it is not visible, texturing and shading of the surface is not necessary. One advantage of this system over the Z-buffer is that non-visible surfaces do not have to be textured or shaded.

25 Texturing or shading of images requires a great deal of processing power. The reduction in processing achieved by the system of the above-mentioned United States Patent is therefore very useful, and can be quite dramatic with certain types of image, particularly those having many
30 overlapping polygons or surfaces.

35 In the system of the United States Patent, the surfaces defining an object are assumed to extend across the whole of the image plane, that is to say they are 'infinite' in extent. Also, each surface is defined as being a forward surface, if it is at the front of the

- 3 -

object and thus faces towards the observer, or a reverse surface if it forms part of the back of the object and thus faces away from the observer. To determine whether any given object is visible at any given pixel, in the 5 ray-casting technique, a comparison is made of the distances from the observation point to (a) the forward surface intersection with the ray which is furthest from the observation point and (b) the reverse surface intersection with the ray which is closest to the 10 observation point. If (a) is greater than (b), then, as illustrated in the Patent, that indicates that the ray does not intersect with that object, and thus that that particular object is not visible at the pixel in the image plane through which the ray passes. With this technique, 15 the edges of the surfaces do not have to be defined or calculated as such; it is sufficient to know the vertices of the object as a whole and to calculate the planes occupied by the surfaces. It will be appreciated that the technique requires that every object is checked for every 20 pixel in the image plane to determine whether or not that object is visible at that location.

As described in that patent, the technique makes it particularly easy to apply shadows to appropriate parts of the image. The system is also able to deal with 25 transparency which can take various forms.

The technique has advantages over the Z-buffer system, but nevertheless processing requirements can still be a constraint. It is proposed in the Patent to improve performance by subdividing the image plane or screen into 30 a plurality of sub-regions or 'tiles'. The tiles are conveniently rectangular (including square). Then for each tile, those objects having surfaces which could fall within the tile are first determined, and only those objects within the tile are processed, thus decreasing the 35 number of surfaces to be processed. The determination of

- 4 -

which objects could contribute to each tile may be achieved by surrounding the object with a bounding volume, namely a cuboid which fully contains the object, and comparing the tile area with the bounding volume. To do 5 this, all the bounding volumes are projected onto the image plane and are tested against the corners of the tiles. Those objects which have bounding surfaces which are completely outside the tile are discarded for that tile. Thus the number of surfaces which need to be 10 processed per pixel within a tile becomes less, and hence the total time to render an image is reduced, since the total processing time for all the tiles will be reduced.

The United States Patent describes the use of tiles of variable size. This reflects the fact that objects are 15 not normally evenly distributed over the entire screen. As shown in Figure 1 of the drawings of the present application, three tiles 10 are 10 pixels square to accommodate three particular objects, and four tiles 12 are 5 pixels square to accommodate four smaller objects. 20 The image portions for the several tiles are processed in pipeline fashion in a processing system as described in the patent.

We have appreciated that even despite all these features, processing power can still be a constraint, but 25 can be improved by use of the present invention.

Summary of the Invention

The invention in its various aspects is defined in the appended claims to which reference should now be made. Preferred features of the invention are set forth in the 30 appendant claims.

A preferred embodiment of the invention is described in more detail below with reference to the drawings. Briefly, this preferred embodiment of the invention takes the form of image processing apparatus for rendering (i.e.

- 5 -

coloring, texturing or shading) an image includes a tiling device which divides the image into sub-regions or tiles. Two rendering devices are provided, and the tiles are allocated so that some are processed by one rendering 5 device and some by the other. Polygons representing surfaces of objects to be displayed are tested against the tiles. If the surface falls into one sub-region only, the data is sent to one rendering device only. On the other hand, if the surface falls into two sub-regions being 10 handled by the different rendering devices, then the data is sent to both rendering devices. The result is that a substantial proportion of the data need only be supplied to and processed by one rendering device, thereby speeding the operation of the apparatus. The outputs of the two 15 rendering devices are subsequently combined by tile interleaving and image display circuitry.

Brief Description of the Drawings

The invention will now be described in more detail by way of example with reference to the accompanying 20 drawings, in which:

Figure 1 illustrates a portion of an image plane containing several objects and having tiles of variable sizes, taken from U.S. Patent 5,729,672;

25 **Figure 2** illustrates a portion of the screen containing tiles in accordance with an embodiment of the present invention;

Figure 3 is a diagram showing a bounding box around an object;

30 **Figure 4** is a flow chart of the procedure used to determine the co-ordinates of a bounding box around an object;

Figure 5 illustrates a portion of the screen containing tiles in a modification of the embodiment of Figure 1; and

- 6 -

Figure 6 is a block schematic diagram of hardware used in preferred image processing apparatus embodying the invention.

Detailed Description of the Preferred Embodiments

5 A method of and apparatus for image processing, more particularly for the real-time texturing or shading of three-dimensional (3D) images, will now be described with reference to Figure 2 *et seq.* of the drawings.

10 The introduction above describes a "bounding volume" technique in which objects which have bounding volumes which are completely outside a tile are discarded for that tile. In the following description a specific 15 implementation of this principle is used, namely it is assumed that objects are projected onto the screen or image plane, and their bounding surfaces as seen on that plane are compared with the tiles.

20 Referring first to Figure 2, there is shown a portion 22 of a screen which contains twelve sub-regions or tiles 22 as shown. Each tile is typically 32 pixels or 64 pixels square, that is, using a conventional raster scan, 25 32 or 64 pixels long by 32 or 64 lines high. It will be seen that two objects are diagrammatically shown on this Figure, namely a house 24 and a bicycle 26. Each of these objects extends, for the sake of illustration, over two of the tiles 22. In the case of the house 24 these are one above the other in the image, and in the case of the bicycle they are side by side. Either way the operation is the same.

30 The tiles are divided into two groups of tiles. As shown, the tiles are split into two groups in checkerboard fashion. As shown in Figure 2, alternate tiles are shaded light or dark on the figure, with each light tile 22A being surrounded by four dark tiles and each dark tile 22B

- 7 -

being surrounded by four light tiles. The light and dark tiles thus form diagonals across the screen.

In accordance with this invention, the rendering of the objects, that is to say the texturing and shading of 5 the objects, which has hitherto been handled by a single processor, is split between two processors which may be referred to as processor A and processor B. Each group of tiles is associated with a respective one of the processors A and B. That is, all the light tiles 22A are 10 associated with processor A and all the dark tiles 22B, as shown in Figure 2, are associated with processor B. All the processing of surfaces which are seen in a light tile is undertaken by processor A and all the processing of surfaces which are seen in a dark tile is undertaken by 15 processor B.

Complex objects can be seen to be made up of a group of several smaller polygons. For example, the house 24 is made up of a triangle for the roof 28, a triangle and a square for the chimney 30, and rectangles for the main 20 body 32 of the house, the windows 34, and the door 36.

Those polygons which make up the roof 28, the chimney 30, and the upstairs one of the windows 34 lie entirely within a light tile 22A, and so need only be sent to processor or device A. The door 36 and the downstairs 25 window 34 lie entirely within a dark tile 22B, and so need only be sent to processor or device B. However, the main body 32 of the house overlaps two tiles. Accordingly it must be sent to both processors or devices as it affects the display in both a light tile and a dark tile.

30 Similarly for the bicycle 26, the polygons which make up the front wheel, the front forks and the handlebars are sent to device A only; the rear wheel, the rear forks and the saddle are sent to device B only; while the frame of the bicycle is sent to both devices. Larger objects may 35 extend over three or more tiles, but the processing

- 8 -

applied and its effects are the same as with objects that overlap just two tiles.

The extent of each of the polygons which make up the complex object can be achieved by drawing a rectangle to 5 enclose the entire polygon, and then testing the rectangle against the co-ordinates of the tiles of the screen to determine the extent of the polygon. This is illustrated in Figure 3, which shows an arbitrary polygon 40. This polygon is shown as a re-entrant polygon and may 10 conveniently be broken into non-re-entrant polygons for processing, if desired, but the principle of establishing the bounding box is the same. Polygon data is supplied to the rendering device by giving the co-ordinates of each vertex of the polygon. The co-ordinates are given in a 15 Cartesian system which has three orthogonal axes, X, Y and Z. The final display screen is assumed to be in the plane of the X-Y axes, while the Z axis represents the depth of the object, as is conventional. The procedure used to determine the minimum values of x and y which define the 20 corners 44 of the bounding box 42 is illustrated in Figure 4.

In the procedure 50 of Figure 4, in a first step 52 the input values x_{in} , y_{in} for the first vertex processed are initially assumed to be the desired values x_{min} , x_{max} , y_{min} , 25 y_{max} . The next vertex is then processed in step 54. The new input values x_{in} , y_{in} are compared with the stored values for x_{min} , x_{max} , y_{min} , y_{max} . If for either x or y the new input value is less than the stored minimum value then it replaces the minimum value, and if the new input value 30 exceeds the maximum value it replaces the maximum value. The procedure then moves to step 56 where a check is made to see whether the last vertex defining the polygon has been processed. If not, the procedure returns to step 54 to process the next vertex. If it has, then the 35 co-ordinates of the bounding box 42 have now been

- 9 -

determined, and the procedure moves to step 58 where a determination is made to see whether the bounding box overlaps the screen tile which is being processed. This determination also is made by simple comparison of the XY 5 co-ordinates of the corners 44 of the bounding box 42 with the XY co-ordinates of the corners of the tile. The tile size in terms of the number of pixels is preferably chosen to be a power of two, e.g. 32 or 64, which has the consequence that the screen tile test 58 reduces to a 10 number of simple binary comparisons.

The bounding box procedure described will occasionally cause the system to indicate that a surface falls within a given tile when in fact it does not, but this is not a significant problem.

15 As described, the tiles have been divided between the two processors A and B in a checkerboard pattern. Other methods may however be used to divide up the screen. The optimal method will depend on the image content, and may mean that more tiles are processed by one processor than 20 the other. Figure 5 illustrates a form of division in which a large rectangle 2 contains a small rectangle 1 of shorter height and width. The device B is arranged to process the tiles forming the smaller rectangle 1, and the processor A is arranged to process what is left, that is 25 rectangle 2 with rectangle 1 excluded. Rectangle 1 is thus defined for processor A as an inclusion rectangle and rectangle 2 as an exclusion rectangle. That is, all polygons entirely enclosed in rectangle 1 are sent to device A. Polygons outside rectangle 1 but inside 30 rectangle 2 are sent to device B. Polygons which overlap the two areas are sent to both devices.

The division of the screen is such that a substantial number of surfaces have to be sent only to one processor. Thus the tiles are preferentially substantially square 35 with, say, one side not more than twice or three times the

- 10 -

length in pixels of the other, so that there is a reasonable likelihood that a good proportion of the surfaces will fall only in one tile. In this way the processing is reduced by some surfaces having to be sent 5 to one processor only. Thus, a split in which alternate scan-lines were processed by different processors would not provide any advantage because the number of surfaces which fall only on one scanline is zero, or close to it. Another possible split would be in horizontal bands across 10 the screen, but the bands would have to be sufficiently wide, e.g. the screen as a whole might be split into three or four bands. Then the tile aspect ratio is less than six to one. In any event, there will normally be at least three screen areas, with at least one of the processors 15 being arranged to process at least two discrete separate image areas.

The hardware required for implementation of the embodiment described takes the form of that shown in Figure 6. Figure 6 shows in block schematic form image processing apparatus 60 embodying the invention and comprising a central processing unit (CPU) 62 connected to a main memory 64. A tiling device 66 defines the tiles and communicates with a local memory 68 as well as with the CPU 62. The tiling device effectively has two outputs 25 to which are attached a first texturing or rendering device 70A and a second texturing or rendering device 70B. The outputs of the two rendering devices 68A and 68B are both applied to tile interleaving and image display circuitry 72. For further description of suitable hardware, reference may be made to our above-mentioned 30 United States Patent, the significant point being that there are two rendering devices 70A, 70B connected in parallel between the tiling device 66 and the circuitry 72.

- 11 -

In practice the two outputs of the tiling device 66 may be constituted by a single data bus together with appropriate addressing to identify a required one of the two rendering devices 70A, 70B.

5 The steps involved in the operation of the apparatus of Figure 6 may be described in principle as follows:

1. Objects are generated by the user (programmer). They are defined by their vertices, and by texture codes which indicate the type of texturing required for each surface. This will include the color and other surface effects.
2. A bounding box is generated for each surface, following the method described above.
3. The bounding box is compared against the macro tiling pattern being employed, so as to determine which surfaces fall into which tiles, and hence which of the multiple rendering devices (two in this case) require the data. Surface vertices and texture codes are then stored in the appropriate local memory portion associated with each memory device.
4. At this stage a tile display list for each rendering device is generated, so that when the rendering device operates, it only needs to traverse data for each of its tiles, and not the whole scene display list.
5. For each pixel of the tile the surfaces are sorted by depth, with the nearest surfaces listed first.
6. The ray-casting method of the above-mentioned U.S. Patent is employed to find the front-most opaque surface which is seen at that particular pixel.
7. The thus-located front-most visible surfaces are then rendered, to provide the desired surface color, texture, and shading.
8. The resultant is stored in a display buffer.

- 12 -

In accordance with this invention, some of the above steps are executed by two processors operating in parallel. The operations are distributed between the two processors in the manner described above. That is to say, 5 step 7 for some tiles (the light tiles as referred to above) is achieved by one processor and the same step for the other (dark) tiles is achieved by the other processor. This reduces the time required for processing by a factor less than but approaching a half. The outputs of the two 10 processors are combined for application to the display buffer and subsequent display on the display screen.

It will therefore be seen that the system operates by supplying data defining a group of surfaces representing an object, e.g. the house or the bicycle in Figure 2. The 15 display is subdivided into a large number of tiles, and a determination is made as to which surfaces fall into which tiles. The data is then applied to the two rendering devices in dependence upon which tile the various surfaces fall into. The data of some surfaces will be sent to one 20 rendering device only and the data of other surfaces will be sent to both rendering devices. More particularly, when the surface falls into one tile only, e.g. the roof or the door in Figure 2, the data need only be sent to one rendering or texturing device. When the surface falls 25 into two tiles being handled by the different rendering devices, the surface data must be sent to both rendering devices.

It may be that the surface falls into two tiles which are both handled by the same rendering device or 30 processor, in which case, again, the data need be sent only to one rendering device. This is unlikely with the tile arrangement of Figure 2, but could more easily happen with other arrangements.

The embodiments of the invention illustrated assume 35 that there are two processors instead of the usual one to

- 13 -

provide the rendering (texturing or shading for instance).
However, the invention is not limited to the use of two
devices; more than two may be used if desired, in which
case the screen is split up into an appropriate larger
5 number of regions each comprising a respective group of
tiles.

The embodiments illustrated have the advantage that,
assuming normal images are being processed, the processing
time is reduced, due to the fact that the processors can
10 operate simultaneously in parallel, each processor needing
to process only some of the surfaces and not all the
surfaces in the image.

It will be appreciated that many other modifications
may be made to the system described and illustrated purely
15 by way of example.

- 14 -

CLAIMS

1. Image processing apparatus, comprising:
 supply means (62,64) for supplying data defining a group of surfaces representing each object in an image;
 sub-division means (66,68) coupled to the supply means for sub-dividing the display into a plurality of sub-regions, for determining which surfaces may fall into which sub-regions, and for applying surface data to a plurality of output means in dependence upon which sub-regions the surfaces fall;
 a corresponding plurality of rendering devices (70A,70B), each coupled to a respective output of the sub-division means; and
 combining means (72) coupled to the plurality of rendering devices to receive and combine the respective outputs of the plurality of rendering devices for display;
 wherein the data of some surfaces is sent to one rendering device only and the data of other surfaces is sent to more than one rendering device, in dependence upon whether the surface falls into one sub-region only or into a plurality of sub-regions being handled by different rendering devices.
2. Apparatus according to claim 1, in which the supply means comprises a CPU (62) and a memory device (64).
- 25 3. Apparatus according to claim 1, in which the sub-division means (66) has associated memory means (68).
4. Apparatus according to claim 1, in which the sub-division means (66) determines which surfaces may fall into which sub-region by determining a rectangular bounding volume for each surface, and determining which sub-regions the bounding volumes fall into.

- 15 -

5. Apparatus according to claim 1, in which the subdivision means (66) uses a ray-casting method to determine which surfaces may fall into which sub-regions in an image plane.
- 5 6. Apparatus according to claim 1, in which the surfaces are sorted by depth before rendering by the rendering means (70A, 70B).
7. Apparatus according to claim 1, in which the surfaces are polygons.
- 10 8. Apparatus according to claim 1, in which there are two and only two rendering devices (70A, 70B).
9. Apparatus according to claim 8, in which the sub-regions are associated with the rendering devices (70A, 70B) in a checkerboard pattern.
- 15 10. Apparatus according to claim 1, in which the sub-regions consist of respective bands across the image.
11. Apparatus according to claim 1, in which the sub-regions are rectangular with an aspect ratio of less than one in six.
- 20 12. Apparatus according to claim 1, in which the sub-regions are rectangular with an aspect ratio of less than one in three.
13. Apparatus according to claim 1, in which at least some of the sub-regions are substantially square.
- 25 14. Apparatus according to claim 1, in which at least one sub-region is defined by inclusion in a first rectangle

- 16 -

and exclusion from a second rectangle contained within the first rectangle.

15. Apparatus according to claim 1, in which at least some of the sub-regions are from 32 to 64 pixels in width

5 and height.

1 / 4

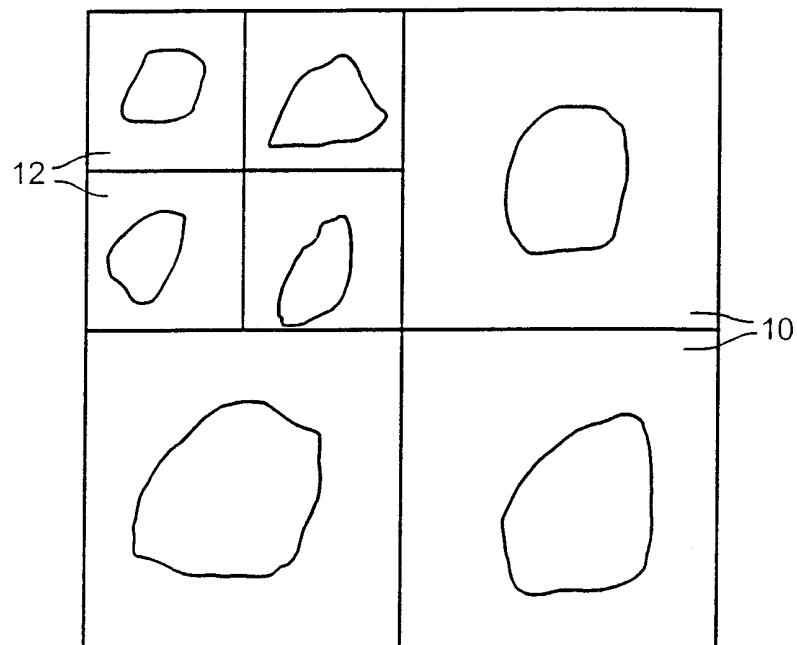


FIG. 1
PRIOR ART

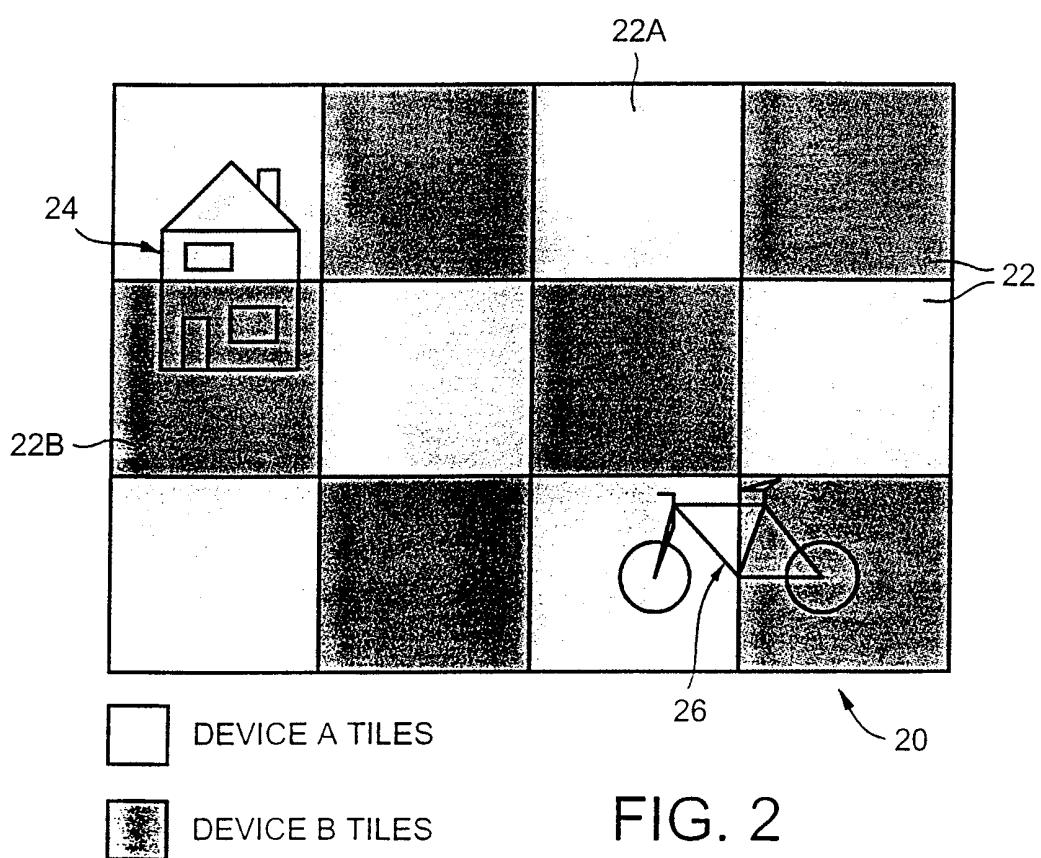


FIG. 2

2 / 4

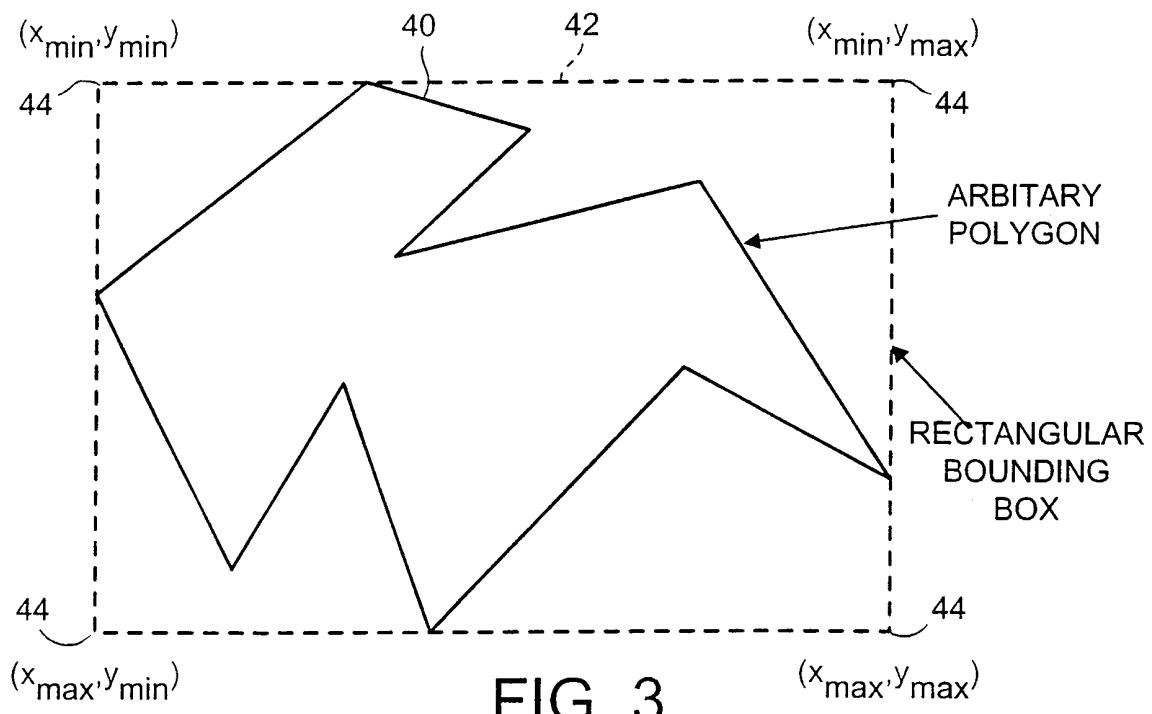
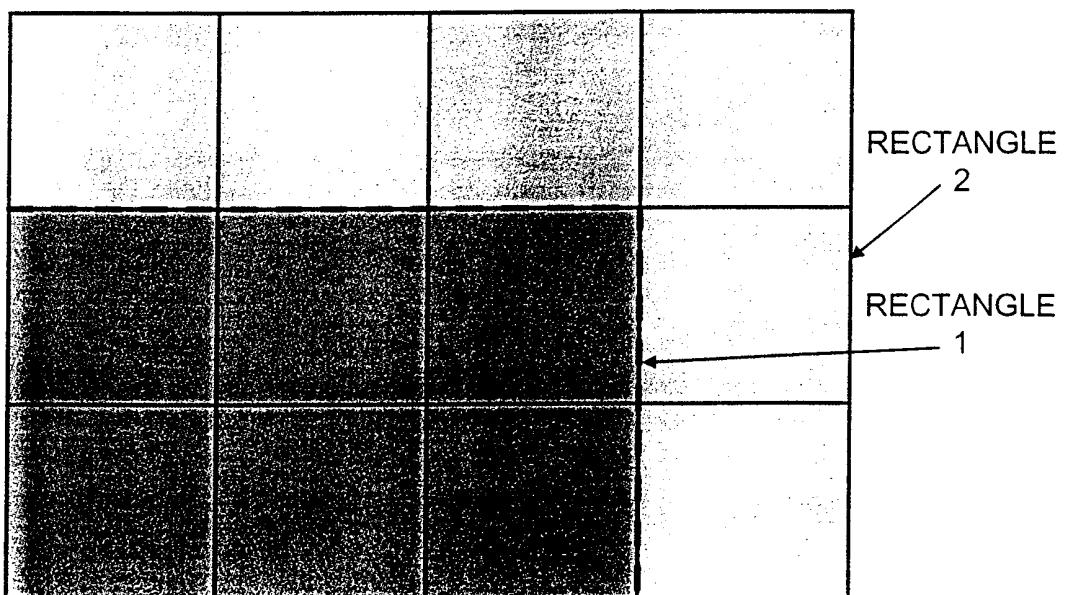


FIG. 3



DEVICE A TILES (INCLUSION R1, EXCLUSION R2)



DEVICE B TILES (INCLUSION R1)

FIG. 5

3 / 4

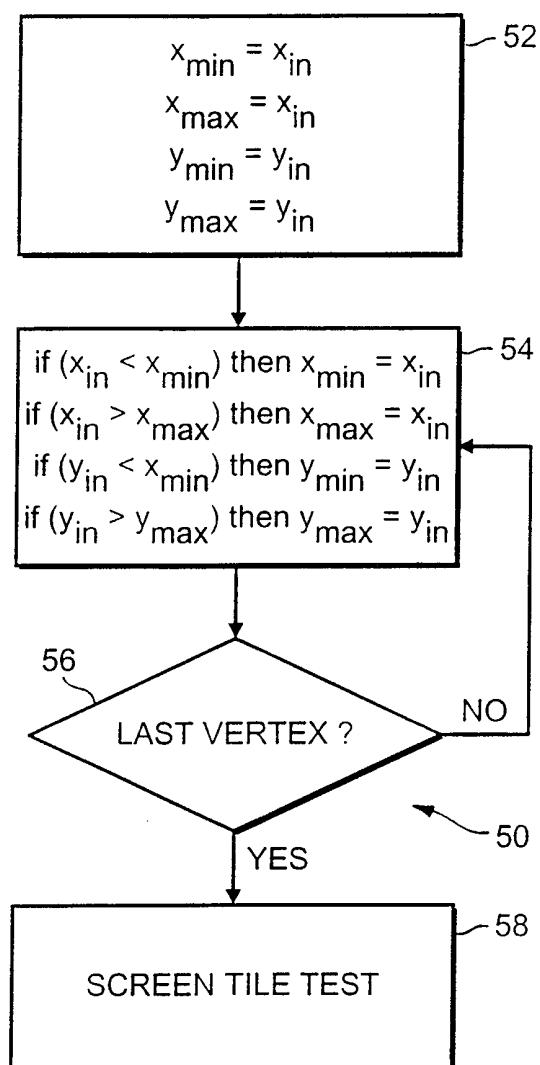


FIG. 4

4 / 4

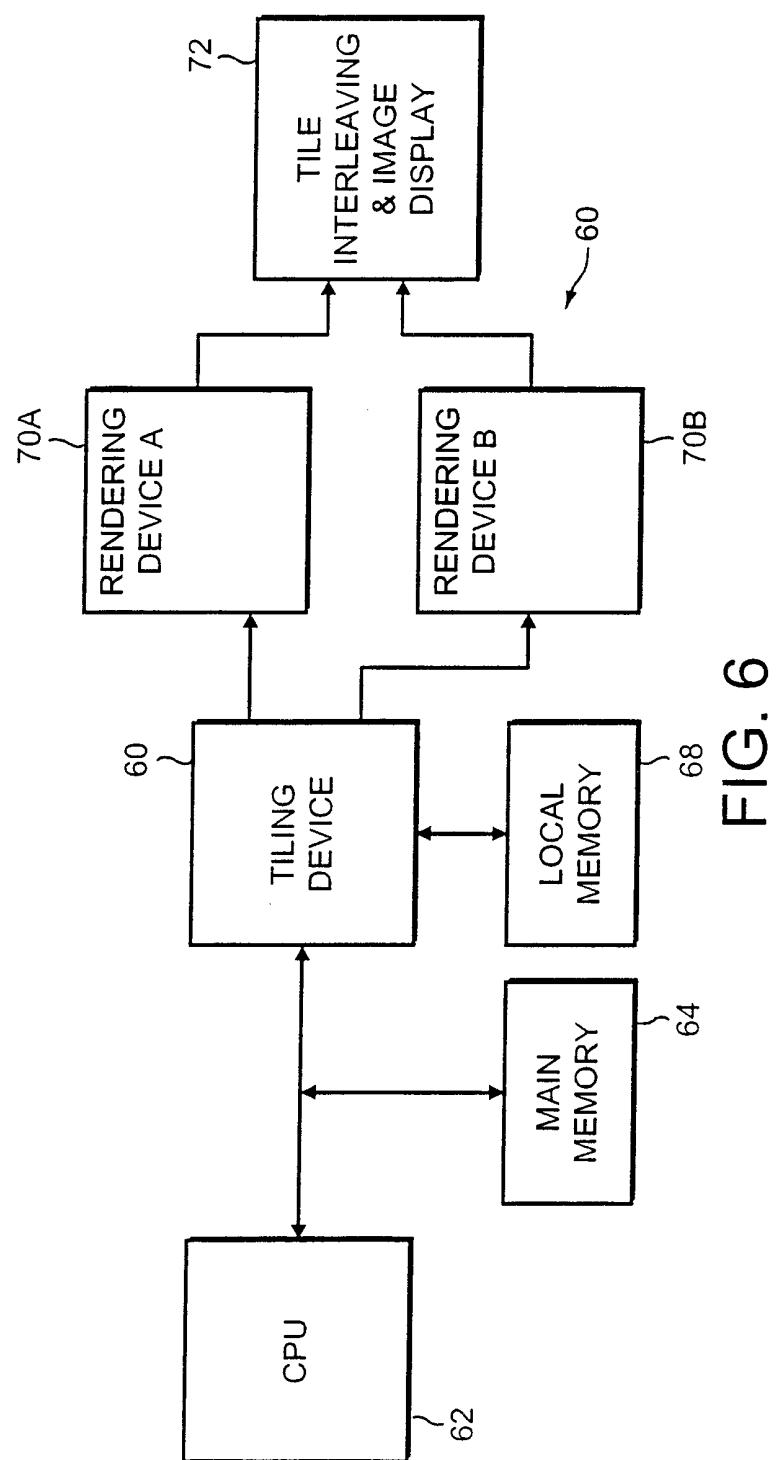


FIG. 6

INTERNATIONAL SEARCH REPORT

Intern. Appl. No

PCT/GB 99/03716

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 G06T15/00

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7 G06T

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the International search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	<p>KAWAI T ET AL: "EVALUATION OF PARELLE RAY TRACING ALGORITHM USING SCREEN SPACE SUBDIVISION FOR IMAGE GENERATION SYSTEM MAGG" SYSTEMS & COMPUTERS IN JAPAN, US, SCRIPTA TECHNICA JOURNALS. NEW YORK, vol. 25, no. 11, October 1994 (1994-10), page 78-86 XP000499051 ISSN: 0882-1666 page 78, left-hand column, line 15 - line 20 page 79, right-hand column, line 9 - line 11 page 82, left-hand column, line 4 - line 8 ----- -----</p>	1-5, 7, 9, 11

Further documents are listed in the continuation of box C.

Patent family members are listed in annex.

* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the International filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- "&" document member of the same patent family

Date of the actual completion of the International search

21 February 2000

Date of mailing of the International search report

28/02/2000

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3018

Authorized officer

Burgaud, C

INTERNATIONAL SEARCH REPORT

Internat'l Application No

PCT/GB 99/03716

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	<p>COPPEN D ET AL: "A DISTRIBUTED FRAME BUFFER FOR RAPID DYNAMIC CHANGES TO 3D SCENES" COMPUTERS AND GRAPHICS, GB, PERGAMON PRESS LTD. OXFORD, vol. 19, no. 2, March 1995 (1995-03), page 247-250 XP000546629 ISSN: 0097-8493 page 247, left-hand column, line 17 - line 19 page 247, right-hand column, line 1 - line 29 page 248, left-hand column, line 4 - line 15</p> <p>—</p>	1-4,7,9, 11