



(51) International Patent Classification:

G06T 11/00 (2006.01) G06T 15/50 (2006.01)
G06T 15/04 (2011.01)

(21) International Application Number:

PCT/GB2016/052432

(22) International Filing Date:

5 August 2016 (05.08.2016)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

1518908.7 26 October 2015 (26.10.2015) GB

(71) Applicant: ARM LIMITED [GB/GB]; 110 Fulbourn Road, Cambridge Cambridgeshire CB1 9NJ (GB).

(72) Inventor: PATEL, Nikunj; ARM Limited, 110 Fulbourn Road, Cambridge Cambridgeshire CB1 9NJ (GB).

(74) Agent: DEHNS; St. Bride's House, 10 Salisbury Square, London Greater London EC4Y 8JD (GB).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM,

AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.


(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

— with international search report (Art. 21(3))

(54) Title: GRAPHICS PROCESSING SYSTEMS

Fig. 2



(57) Abstract: In a graphics processing system, a first bounding volume (20) representative of the volume of all or part of a scene to be rendered is defined. Then, when rendering a scene for output, an intersection position (27) for a vector (24) with the first bounding volume (20) is determined by determining the position (51) intersected by the vector (24) on a second bounding volume (50) associated with the first bounding volume (20), using the determined position (51) to determine information indicative of an intersection position (27) for the first bounding volume (20) and the vector (24), and using the determined information to determine an intersection position (27) for the vector (24) with the first bounding volume (20).

Graphics Processing Systems

5

BACKGROUND

10 The technology described herein relates to graphics processing systems and in particular to methods of and apparatus for determining the position on a volume intersected by a vector when rendering images for display.

When rendering images, such as output frames, for display in graphics processing systems, it can often be necessary to determine the position on a volume intersected by a vector. For example, in environment mapping (e.g. reflection mapping) techniques, the position on a bounding volume intersected by a (e.g. reflected) vector is determined. The intersection position is then used to determine a sampling vector to be used to sample a graphics texture that represents the surface of the bounding volume.

15 In these techniques, the bounding volume is used to approximate the environment (e.g. room) of the scene being rendered. Typically, a relatively simple bounding volume, such as a bounding box (a cube) or a bounding sphere, is used so that the intersection point can be analytically determined in real-time (at run time) in a relatively simple and efficient manner.

20 While these techniques can be effective where the environment of the scene being rendered corresponds closely to the, e.g. cubic or spherical, bounding volume being used (e.g. to a box or sphere), it can often be the case that the bounding volume does not match exactly to the surface of the scene being defined. This can cause noticeable errors in the rendered image, e.g. in the reflections, generated using the environment mapping technique.

30 The Applicants believe therefore that there remains scope for improved graphics processing techniques where the position on a volume intersected by a vector is determined.

35

BRIEF DESCRIPTION OF THE DRAWINGS

A number of embodiments of the technology described herein will now be described by way of example only and with reference to the accompanying drawings, in which:

Figure 1 shows schematically a graphics processing pipeline that can be operated in the manner of the technology described herein;

Figure 2 shows schematically the operation of an embodiment of the technology described herein;

Figure 3 shows schematically the generation of a texture in an embodiment of the technology described herein;

Figure 4 shows an exemplary texture used in an embodiment of the technology described herein;

Figure 5 shows schematically the generation of a texture in an embodiment of the technology described herein; and

Figure 6 illustrates the operation of an embodiment of the technology described herein.

Like reference numerals are used for like components where appropriate in the drawings.

DETAILED DESCRIPTION

A first embodiment of the technology described herein comprises a method of operating a graphics processing system when rendering a scene for output, in which a first bounding volume representative of the volume of all or part of the scene to be rendered is defined; the method comprising:

determining an intersection position for a vector with the first bounding volume by:

determining the position intersected by the vector on a second

bounding volume associated with the first bounding volume;

using the determined position to determine information indicative of an intersection position for the first bounding volume and the vector; and

using the determined information to determine an intersection position for the vector with the first bounding volume.

- 3 -

A second embodiment of the technology described herein comprises a graphics processing unit comprising:

processing circuitry configured to determine an intersection position for a vector with a first bounding volume representative of the volume of all or part of a scene to be rendered:

5

determining the position intersected by the vector on a second bounding volume associated with the first bounding volume; and

using the determined position to determine information indicative of an intersection position for the first bounding volume and the vector; and

10

using the determined information to determine an intersection position for the vector with the first bounding volume.

The technology described herein is directed to a method of and apparatus for determining an intersection position for a bounding volume and a vector. In the technology described herein, to determine an intersection position for a vector with a first bounding volume, instead of “directly” determining the position intersected by the vector on the first bounding volume, the position intersected by the vector on a second bounding volume associated with the first bounding volume is firstly determined. The so-determined intersection position is used to determine information indicative of an intersection position for the first bounding volume and the vector, and the determined information is used to determine the intersection position for the vector with the first bounding volume.

15

20

As will be discussed further below, the Applicants have recognised that this arrangement can be used to determine an intersection position for an arbitrary volume and a vector in a particularly convenient and efficient manner (e.g. when compared with “directly” determining the position intersected by the vector on the volume in question). Moreover, this means that the technology described herein can allow the first bounding volume that represents the volume of all or part of the scene, e.g. in environment mapping techniques, to take a relatively complex shape, e.g. so as to correspond more closely to the (surface of the) environment of the scene being rendered (e.g. where the environment is not simply a box or a sphere), without significantly increasing the processing required for determining the intersection position for that bounding volume and the vector. This can accordingly simplify the processing required for determining the position on a bounding volume intersected by a vector, e.g. for use in determining a sampling vector to be used to

25

30

sample a graphics texture that represents the surface of the bounding volume in environment mapping techniques.

5 The technology described herein can accordingly significantly reduce the processing required as compared to conventional graphics processing techniques and/or can provide enhanced rendering quality. This can be particularly useful for rendering images on, e.g., a mobile device that may have limited processing capacity.

10 The first bounding volume that is defined and used in the technology described herein may represent the entire scene that is being rendered (and in one embodiment this is the case), but it would also be possible for the bounding volume to represent only a part of the scene. The first bounding volume may, for example, represent a room or some other environment of the scene to be rendered.

The first bounding volume can be defined in any desired and suitable manner. It is in an embodiment defined in world space.

15 The first bounding volume can take any suitable and desired form. The first bounding volume in an embodiment comprises a convex volume, i.e. such that for any two points within the bounding volume, every point on the straight line segment that joins the two points is also within the first bounding volume. However, it would also be possible for the first bounding volume to comprise a shape that is not
20 "strictly" convex in this sense. For example, the first bounding volume could comprise a (substantially convex) volume having one or more holes. As will be described in more detail below, the one or more holes in the first bounding volume could be used, e.g. to represent windows in the environment (e.g. room) of the scene being rendered.

25 The first bounding volume could be in the form of a bounding box (a cube), or a bounding sphere. However, the first bounding volume could also be (and in an embodiment is) in the form of any other shape (i.e. other than (not) a box or a sphere) and/or a more complex shape (and indeed, as discussed above the technology described herein is particularly advantageous in that the first bounding
30 volume can take any arbitrary shape). The first bounding volume in an embodiment corresponds closely to the environment of (to a surface of) the scene being rendered.

The second bounding volume that is defined and used in the technology described herein should be associated with the first bounding volume. The second

bounding volume can be defined in any desired and suitable manner. It is in an embodiment defined in world space.

The second bounding volume should be (and in an embodiment is) configured such that the vector will intersect with the second bounding volume (thus, the second bounding volume should be (and in an embodiment is) configured such that the origin of the vector is within the second bounding volume). The second bounding volume in an embodiment at least partially overlaps the first bounding volume. The second bounding volume is in an embodiment centred on the first bounding volume. The second bounding volume in an embodiment encompasses (bounds) the first bounding volume (i.e. such that the origin of the vector will be within the second bounding volume where the origin of the vector is within the first bounding volume), although it would also be possible for the second bounding volume to only partially encompass (partially bound) the first bounding volume, to be bounded by the first bounding volume and/or to only be partially bounded by the first bounding volume.

The second bounding volume can take any suitable and desired form. The second bounding volume is in an embodiment representative of the first bounding volume. In an embodiment, the second bounding volume has a form that is relatively simple when compared with the form of the first bounding volume. For example, the second bounding volume may be a simplified version of the first bounding volume. This then means that the position intersected by the vector on the second bounding volume is in an embodiment analytically determinable in a relatively simple and efficient manner, e.g. when compared with the processing that would otherwise be required to analytically determine the position intersected by the vector on the (relatively more complex) first bounding volume.

In an embodiment, the second bounding volume has the form of a box (a cube) or a sphere. Other arrangements for the second bounding volume would, of course, be possible.

The vector for which the intersection position is determined in the technology described herein may comprise any suitable such vector for which it is desired to determine an intersection position with respect to the first bounding volume. The vector should be (and in an embodiment is) defined by a direction, optionally together with a magnitude (length) (which may correspond to unity), a starting point (origin) and/or an end point. The starting point (origin) of the vector should be (and in an embodiment is) within the first (and second) bounding volume.

As will be described in more detail below, the vector is in an embodiment associated with a sampling point on or within the first bounding volume for which it is desired to determine a property (e.g. colour, transparency, etc.) in the rendering process (e.g. by sampling a graphics texture that represents the surface of the first bounding volume using a sampling vector determined from the intersection position for the first bounding volume and the vector). In these embodiments, the vector in an embodiment starts (or ends) at the sampling point.

As will be described in more detail below, although it is possible that the intersection position for the first bounding volume and the vector that is determined in the technology described herein will be (precisely) equivalent to the position on the first bounding volume intersected by the vector, it will more generally be the case that the intersection position for the first bounding volume and the vector that is determined in the technology described herein will be an approximation to the position on the first bounding volume intersected by the vector.

In contrast, the position on the second bounding volume intersected by the vector is in an embodiment analytically determined, e.g. using information defining the vector and the second bounding volume, although it would also be possible to determine an approximation to the position on the second bounding volume intersected by the vector, if desired.

The position on the second bounding volume intersected by the vector can be determined in any suitable and desired manner. In an embodiment, this is done in world space.

The determined intersection position on the second bounding volume (that is used in the later stages of the process of the technology described herein, e.g. in determining the intersection position for the first bounding volume and the vector) may take the form of a set of coordinates (in world space). However, in an embodiment, the intersection position on the second bounding volume takes the form of (is represented as) (or is used to derive) a direction from a reference position for the second bounding volume (e.g. that the second bounding volume is defined with respect to), e.g., a unit vector and/or a pair of spherical coordinates, etc.

In these embodiments, the reference position for the second bounding volume is in an embodiment within the second bounding volume, i.e. such that each direction from the reference position uniquely defines a position on the surface of the second bounding volume. The reference position for the second bounding

- 7 -

volume (that the second bounding volume is defined with respect to) is in an embodiment at the centre of the second bounding volume, although other arrangements would be possible.

5 These processes of determining the position on the second bounding volume intersected by the vector can be carried out by any desired and suitable stage or component of the graphics processing unit. In an embodiment, the position on the second bounding volume intersected by the vector is determined by a fragment shading stage of the graphics processing unit, in an embodiment by that stage executing an appropriate fragment shading program (fragment shader) that
10 can calculate the intersection point.

Other arrangements, such as this determination being done by a vertex shading stage of the graphics processing unit would also be possible, if desired.

As discussed above, in the technology described herein, the determined position on the second bounding volume intersected by the vector is used to
15 determine information indicative of an intersection position for the first bounding volume and the vector. This may be done in any suitable and desired manner.

In an embodiment, the position on the second bounding volume intersected by the vector is used together with the direction of the vector to determine the information. As will be appreciated by those having skill in the art, the position on
20 the first bounding volume intersected by the vector will be uniquely defined by these two pieces of information.

In these embodiments, the direction of the vector may take the form of (and/or may be used to derive) a unit vector and/or pair of spherical coordinates, etc.

25 In an embodiment, the position on the second bounding volume intersected by the vector is used (in an embodiment together with the direction of the vector) to look up information indicative of an intersection position for the first bounding volume and the vector, in an embodiment from pre-determined information. As will be appreciated by those having skill in the art, looking up this information
30 represents a particularly convenient and efficient technique for determining the information indicative of the intersection position for the first bounding volume and the vector, e.g. when compared with calculating this information in real-time (at run time).

Thus, in an embodiment, using the determined position to determine
35 information indicative of the intersection position for the first bounding volume and

the vector comprises using the determined position together with the direction of the vector to look up information indicative of an intersection position for the first bounding volume and the vector from pre-determined information.

5 In these embodiments, the pre-determined information in an embodiment comprises information indicative of a position on the first bounding volume intersected by the vector for each of plural, in an embodiment selected, possible combinations of the position on the second bounding volume intersected by a vector and the direction of the vector.

10 The pre-determined information in an embodiment comprises information indicative of a position on the first bounding volume intersected by the vector for less than all possible combinations of the position on the second bounding volume intersected by a vector and the direction of the vector, and so the intersection position for the first bounding volume and the vector that is determined in the technology described herein will only be (precisely) equivalent to the position on the
15 first bounding volume intersected by the vector in a limited number of cases, and will more generally be an approximation to the position on the first bounding volume intersected by the vector.

In one embodiment, the information indicative of an intersection position for the first bounding volume and the vector comprises position information, e.g. the
20 (coordinates of the) intersection position for the first bounding volume and the vector. Correspondingly, according to an embodiment, the pre-determined information comprises position information, e.g. comprising the (coordinates of the) position on the first bounding volume intersected by a vector for each of plural, in an embodiment selected, combinations of the position on the second bounding volume
25 intersected by a vector and the direction of the vector.

Thus, according to an embodiment, using the determined position to determine information indicative of an intersection position for the first bounding volume and the vector, and using the determined information to determine an intersection position for the vector with the first bounding volume comprises:

30 using the determined position (in an embodiment together with the direction of the vector) to determine (to look up) an intersection position for the vector with the first bounding from the pre-determined information.

In an embodiment the coordinates of the intersection position are looked up. In this case, the pre-determined information will indicate the intersection position for
35 the vector with the first bounding volume directly (i.e. in itself).

In another embodiment, the information indicative of an intersection position for the first bounding volume and the vector is information that does not directly indicate the intersection position in itself (on its own) but rather is information that can be (and that is to be) used, e.g. and in an embodiment with other information, to derive and/or determine an intersection position for the first bounding volume and the vector.

In this case, the information indicative of an intersection position for the first bounding volume and the vector may comprise any suitable and desired information that can be used to determine an intersection position for the first bounding volume and the vector. In a particularly such embodiment, the information comprises distance information indicative of the distance from the determined position intersected on the second bounding volume by the vector to an intersection position for the first bounding volume and the vector.

Thus, according to an embodiment, the pre-determined information comprises distance information, e.g. that is indicative of the distance from the position on the second bounding volume intersected by a vector to a position on the first bounding volume intersected by the vector for each of plural, in an embodiment selected, combinations of the position on the second bounding volume intersected by a vector and the direction of the vector.

In these embodiments, the intersection position for the vector with the first bounding volume is in an embodiment determined using the distance information. Thus, according to an embodiment, using the determined position to determine information indicative of an intersection position for the first bounding volume and the vector, and using the determined information to determine an intersection position for the vector with the first bounding volume comprises:

using the determined position to determine a distance value indicative of the distance from the determined position intersected on the second bounding volume by the vector to an intersection position for the first bounding volume and the vector; and

using the determined distance value to determine an intersection position for the vector with the first bounding volume.

In these embodiments, the determined distance value may be used to determine the intersection position for the first bounding volume and the vector in any suitable and desired manner.

In an embodiment, the distance value is used together with the direction of the vector and the position on the second bounding volume intersected by the vector to determine the intersection position for the first bounding volume and the vector. In an embodiment, a distance vector is generated using the distance value together with the direction of the vector, e.g. by multiplying the (normalized, unit) vector by the distance value. This distance vector is in an embodiment then applied to (i.e. added to or subtracted from) the position on the second bounding volume intersected by the vector to determine the intersection position for the first bounding volume and the vector.

Other arrangements would, of course, be possible.

It would also be possible for the pre-determined information to comprise at least some information that is other than (i.e. not) indicative of an intersection position. For example, where there is no intersection for a particular combination of the position on the second bounding volume intersected by a vector and the direction of the vector, and/or where the first bounding volume comprises a volume having one or more holes that are used e.g. to represent windows in the environment (e.g. room) of the scene being rendered (as discussed above), the corresponding values of the pre-determined information (that would be looked up using the position on the second bounding volume intersected by the vector and the direction of the vector) may comprise one or more particular, in an embodiment selected, in an embodiment predetermined, "special", values, such as positive infinity when using floating point values, that are predefined as indicating that the vector will not intersect with the first volume at that position. In this case, where such a "special" value is returned (looked up), then rather than determining an intersection position in the manner of the embodiments, a particular, in an embodiment selected, in an embodiment predefined operation, such as applying a default colour, is in an embodiment performed instead.

The pre-determined information may take any suitable and desired form. In one embodiment, the pre-determined information comprises an array of pre-determined values, where at least some values (and in an embodiment each pre-determined value) in the array in an embodiment comprises a (e.g. position or distance) value indicative of the position on the first bounding volume intersected by a vector for a particular combination of the position on the second bounding volume intersected by the vector and the direction of the vector.

Each value may be, for example, in the form of a floating point number, and may have a certain precision such as a 16-bit precision, etc. Other arrangements would, of course, be possible.

5 The array of pre-determined values can be a 2D array, a 3D array or a 4D array, etc. For example, where the scene is a 3D scene, such that the position on the second bounding volume intersected by a vector and the direction of the vector can each be defined by a 2D direction (e.g. a pair of spherical coordinates), the array in an embodiment comprises a 4D array. Where the scene is a 2D scene, such that the position on the second bounding volume intersected by the vector and
10 the direction of the vector can each be defined by a 1D direction (e.g. an angle), the array in an embodiment comprises a 2D array.

It would also be possible, where the scene is a 3D scene, for the array to comprise a 3D array. In this case, the 4D information could be, for example, “flattened” into a 3D array. Thus, in an embodiment, the pre-determined
15 information comprises a 3D array of 4D pre-determined information. This would then allow the array to be processed with “standard” 3D techniques (as will be described further below).

The array of pre-determined values can have any suitable and desired resolution. As will be appreciated by those having skill in the art, using a higher
20 resolution array will result in more accurate results, but at the expense of requiring more resources (and vice versa). A 4D array could have a resolution of, for example, 32x32x16x16 data elements, e.g. corresponding to 32x32 possible intersection positions on the second bounding volume and 16x16 possible vector directions. A corresponding 3D flattened array would have a resolution of
25 32x32x256 data elements. Other resolutions would, of course, be possible.

In one embodiment, the pre-determined information takes the form of a graphics texture, and the information indicative of the intersection position for the first bounding volume and the vector is looked up by sampling the graphics texture. In these embodiments, the texture is in an embodiment sampled using the position
30 intersected by the vector on the second bounding volume and the direction of the vector. For example, where the position intersected by the vector on the second bounding volume and the direction of the vector are represented by respective pairs of spherical coordinates, these spherical coordinates are in an embodiment used to sample the texture.

The texture can be sampled in any desired and suitable manner. For example, appropriate filtering (interpolation) processes, such as bilinear filtering, can be used when sampling the texture, if desired. Where the texture comprises a 2D, 3D or 4D texture, the texture may be sampled by performing a 2D, 3D or 4D texture look up operation as appropriate. Where the texture comprises a 3D flattened 4D texture, the texture may be sampled by performing two 3D texture look up operations, and appropriately interpolating between the two results. This latter technique is particularly useful in arrangements where, for example, 4D texture look ups are not supported.

Other arrangements would, of course, be possible.

These processes for determining an intersection position for the first bounding volume and the vector can again be carried out by any desired and suitable stage or component of the graphics processing unit. In an embodiment they are performed by a fragment shading stage (a fragment shader) of the graphics processing unit, in an embodiment by executing an appropriate fragment shading program. They may also or instead be performed at least in part by an appropriate texture mapping stage of the graphics processing unit, if desired.

Other arrangements, such as this determination being done by a vertex shading stage of the graphics processing unit would also be possible, if desired.

As will be appreciated from the above, in an embodiment of the technology described herein, there will be a set of intersection position indicating information associated with the second bounding volume, and the determined intersection position on the second bounding volume and the intersection position indicating information associated with the second bounding volume will be used to determine an intersection position for the first bounding volume and the vector.

In these embodiments, the pre-determined information (texture) is in an embodiment generated in advance of its requirement, e.g. "offline" (rather than being generated in "real-time" (at run-time) as and when it is needed), and in an embodiment stored in memory.

In an embodiment, the pre-determined information (texture) is compressed before it is stored. This can accordingly reduce storage and bandwidth requirements. Any suitable (texture) compression process can be used for this. For example, where the pre-determined information comprises a 3D texture, "standard" 3D texture compression techniques, such as the Adaptive Scalable Texture Compression (ASTC) scheme (e.g. as described in GB 2491687), can be

used. Since generally only an approximate intersection position is determined, relatively heavy compression can be used for the texture.

The generated pre-determined information (texture) may be stored, e.g. on an appropriate portable storage medium, such as a DVD, or in memory, for future use by a graphics processing unit, when it is desired to use the pre-determined information (texture) when rendering, e.g., an image.

The technology described herein also extends to the generation of textures indicative of the position on a first bounding volume intersected by a vector for plural combinations of the position on a second bounding volume intersected by the vector and the direction of the vector. It may be, for example, that the texture(s) would be generated (and stored) separately, and/or in advance, and then, e.g., provided to a graphics processor for use. The technology described herein extends to these activities, as well as to the use of the textures for determining the position on a volume intersected by a vector.

Thus, another embodiment of the technology described herein comprises a method of generating a texture for use in a graphics processing system when rendering a scene for output in which a first bounding volume representative of the volume of all or part of the scene to be rendered is defined, the method comprising:

generating a graphics texture comprising an array of texture texels by:
setting the texel values in the texture such that each texel stores a value indicative of an intersection position on the first bounding volume for a vector for a particular combination of the position on a second bounding volume associated with the first bounding volume intersected by the vector and the direction of the vector;
and

storing data representing the texel values of the texture.

Another embodiment of the technology described herein comprises an apparatus for generating a texture for use in a graphics processing system when rendering a scene for output, in which a first bounding volume representative of the volume of all or part of the scene to be rendered is defined, the apparatus comprising processing circuitry configured to:

generate a graphics texture comprising an array of texture texels by:
setting the texel values in the texture such that each texel stores a value indicative of an intersection position on the first bounding volume for a vector for a particular combination of the position on a second bounding volume associated with

the first bounding volume intersected by the vector and the direction of the vector;
and

storing data representing the texel values of the texture.

Another embodiment of the technology described herein comprises a texture
5 for use in a graphics processing system when rendering a scene for output, in
which a first bounding volume representative of the volume of all or part of the
scene to be rendered is defined, the texture comprising:

an array of texture texels, in which:

the texel values are set such that each texel stores a value indicative of an
10 intersection position on the first bounding volume for a vector for a particular
combination of the position on a second bounding volume associated with the first
bounding volume intersected by the vector and the direction of the vector.

As will be appreciated by those skilled in the art, these embodiments of the
technology described herein can, and in an embodiment do, include any one or
15 more or all of the optional features of the technology described herein, as
appropriate.

Thus for example, the first bounding volume in an embodiment comprises a
convex volume. The second bounding volume in an embodiment bounds and/or is
bounded by the first bounding volume, and in an embodiment has a form that is
20 relatively simple compared to the form of the first bounding volume. Each stored
value indicative of an intersection position on the first bounding volume in an
embodiment comprises a position value or a distance value, e.g. and in an
embodiment as described above.

In an embodiment, as well as generating the texture, the “second” bounding
25 volume that the texture is to be used with is also generated and stored. In an
embodiment data defining the second bounding volumes is generated and stored in
association with (associated with) the texture. Again, this bounding volume
information can be subject to any desired post-processing operations, such as
compression, if desired.

30 Once it has been determined, the intersection position for the first bounding
volume and the vector can be used as desired.

It would be possible for the intersection position to be output as the “final”
result of the processing operation. However, in an embodiment, the intersection
position is used in further processing operation(s).

In one embodiment, the techniques of the technology described herein are used in environment mapping. In these techniques (and in an embodiment of the technology described herein), for at least one sampling point on or within a bounding volume that represents the volume of all or part of the scene to be rendered (i.e. the first bounding volume), one or more properties are determined for use in representing the scene at that sampling point. The one or more properties are in an embodiment determined by using the intersection position for the (first) bounding volume and a vector from the sampling point to determine a sampling vector to be used to sample a graphics texture that represents the surface of the (first) bounding volume, and then sampling the graphics texture using the sampling vector.

Thus, according to an embodiment the intersection position for the first bounding volume and the vector is used to determine a vector to be used to sample a graphics texture that represents the surface of the first bounding volume in the scene, and the determined vector is in an embodiment (then) used to sample the graphics texture to determine a property to be used to represent the scene at the sampling point.

The one or more determined properties are in an embodiment one or more properties of the scene (environment) that affect the sampling point, e.g. due to reflection, shadow and/or refraction, etc. The properties may comprise, for example, one or more colour and/or transparency values. Thus, in various embodiments, the techniques of the technology described herein may be used in reflection mapping, and/or other environment mapping techniques such as the Applicant's shadow mapping techniques described in GB 1413145.2, and/or the Applicant's refraction mapping techniques described in GB 1503391.3 and GB 1503386.3.

In the reflection mapping technique for example (and in an embodiment), for at least one sampling point on a reflective surface in the scene, a reflected vector is determined using the vector from the viewpoint (camera) position for the scene to the sampling point, in an embodiment together with the normal of the reflective surface at the sampling point, e.g. based on the law of reflection. The intersection of the reflected vector on a bounding volume representative of the volume of all or part of the scene (i.e. the first bounding volume) is then determined, and the intersection position is used to determine a vector to be used to sample a graphics texture that represents the colour of the bounding volume in the scene. The

graphics texture is (then) sampled using the determined vector to determine a colour value indicative of the surface of the bounding volume as reflected at the sampling point and seen from the viewpoint (camera) position. Thus, according to an embodiment the vector for which the intersection position for the first bounding volume is determined comprises a reflected vector (in an embodiment determined as described above), and the intersection position for the first bounding volume and the reflected vector is used to determine a sampling vector to be used to sample a graphics texture that represents the colour of the bounding volume in the scene. The determined sampling vector is in an embodiment (then) used to sample the graphics texture to determine a colour to be used to represent the part of the scene that will be reflected from and visible at the sampling point in question.

The techniques of the technology described herein may also be used to detect collisions, e.g. between particles and an environment, e.g. for in-game effects.

Thus, according to another embodiment the vector for which the intersection position for the first bounding volume is determined comprises a velocity vector for a particle, and the intersection position for the first bounding volume and the velocity vector is used to determine whether the particle will collide with the environment, e.g. within a given time period.

It is believed that the idea of using the techniques of the technology described herein in environment mapping is new and advantageous in its own right.

Thus, a further embodiment of the technology described herein comprises a method of operating a graphics processing system when rendering a scene for output, in which a first bounding volume representative of the volume of all or part of the scene to be rendered is defined; the method comprising:

for at least one sampling point on or within the first bounding volume, determining a property to be used to represent the scene at that sampling point by: determining an intersection position for a vector from the sampling point with the first bounding volume by:

determining the position intersected by the vector on a second bounding volume associated with the first bounding volume;
using the determined position to determine information indicative of an intersection position for the first bounding volume and the vector; and
using the determined information to determine an intersection position for the vector with first bounding volume;

- 17 -

using the determined intersection position for the first bounding volume and the vector to determine a sampling vector to be used to sample a graphics texture that represents the surface of the first bounding volume; and

5 using the determined sampling vector to sample the graphics texture to determine a property to be used to represent the scene at the sampling point.

An embodiment of the technology described herein comprises a graphics processing unit comprising:

10 processing circuitry configured to, when rendering a scene for output for which a first bounding volume representative of the volume of all or part of the scene to be rendered is defined:

for at least one sampling point on or within the first bounding volume, determine a property to be used to represent the scene at that sampling point by:

determining an intersection position for a vector from the sampling point with the first bounding volume by:

15 determining the position intersected by the vector on a second bounding volume associated with the first bounding volume;

using the determined position to determine the information indicative of an intersection position for the first bounding volume and the vector; and

20 using the determined information to determine an intersection position for the vector with the first bounding volume;

using the determined intersection position for the first bounding volume and the vector to determine a sampling vector to be used to sample a graphics texture that represents the surface of the first bounding volume; and

25 using the determined sampling vector to sample the graphics texture to determine a property to be used to represent the scene at the sampling point.

As will be appreciated by those having skill in the art, these embodiments of the technology described herein can, and in an embodiment do, include any one or more or all of the optional features of the technology described herein, as appropriate.

30 Thus, for example, the second bounding volume is in an embodiment as described above, e.g. the second bounding volume in an embodiment bounds and/or is bounded by the first bounding volume, and in an embodiment has a form that is relatively simple compared to the form of the first bounding volume. The information indicative of an intersection position for the first bounding volume and
35 the vector is in an embodiment determined by using the determined position to look

up the information from pre-determined information, e.g. and in an embodiment as described above.

The determined property may comprise a colour, and/or transparency value, etc.

5 The first bounding volume is in an embodiment as described above, e.g. the first bounding volume in an embodiment comprises a convex volume, and may be e.g. relatively complex.

10 In these embodiments, the graphics texture that represents the surface of the first bounding volume can take any suitable and desired form. In an embodiment it is in the form of a texture that is used for environment mapping, such as a cube texture (cube map) or a sphere texture (sphere map). The texture is in an embodiment configured to correspond to the first bounding volume for the scene.

15 Thus, in an embodiment, the texture that represents the surface of the first bounding volume comprises a texture that indicates and stores for points on a surface that surrounds a reference position within the volume that the texture encompasses, one or more properties (e.g. colour and/or transparency values) for each point. The texture is in an embodiment sampled based on a sampling vector (direction) from the reference position for the texture (that the texture is defined with respect to).

20 In other words, the texture in an embodiment stores one or more properties (e.g. colour and/or transparency values) for respective directions from a reference position (point) within the volume that the texture represents, and is sampled by determining a direction from the reference position (point) within the texture to a position on the surface that the texture represents.

25 The reference position (point) within the texture that the texture is defined (and sampled) with respect to is in an embodiment at the centre of the volume that the texture encompasses, but other arrangements would, of course, be possible, if desired.

30 The texture in an embodiment has a resolution that will be suitable for the anticipated size of the first bounding volume that it is to be used with (although this is not essential, and the texture could, e.g., be scaled for use, if required).

 The texture could be a single channel texture, storing a single property (e.g. colour or transparency (alpha) channel) only. Alternatively, the texture could store

plural properties (plural data channels), such as plural colour and/or transparency channels. In an embodiment, the texture is an RGB or RGB α texture.

5 The texture is in an embodiment generated by rendering an image of a surface for the first bounding volume from the viewpoint of the reference position for the texture (from the texture reference position). The generated texture may be stored, e.g. on an appropriate portable storage medium, such as a DVD, or in memory, for future use by a graphics processing unit, when it is desired to use the texture when rendering, e.g., an image.

10 In an embodiment, the texture is stored (encoded) as a set of mipmaps (i.e. where multiple versions of the original texture each having different levels of detail (resolution) are stored for use).

The texture is in an embodiment generated in advance of its requirement, e.g. "offline" (rather than being generated in "real-time" as and when it is needed).

15 In an embodiment, the texture is subject to one or more processing operations after it has been generated (and before it is stored for use), such as, and in an embodiment, one or more filtering processes, such as having one or more convolution filters applied to it. In an embodiment the texture is subjected to one or more of: blurring; brightness processing; contrast processing (e.g. enhancement); sharpening, etc.. In an embodiment, the texture is subjected to one or more non-
20 real-time optimisations.

In an embodiment, the texture is also compressed before it is stored. Any suitable texture compression process can be used for this.

25 As discussed above, the texture is in an embodiment generated by rendering an image that will represent the surface of the first bounding volume from the viewpoint of the reference position for the texture. This is in an embodiment done by sampling respective positions on the surface (the image representing the surface) for respective positions on the first bounding volume (in the relevant direction from the reference position for the texture). In this process, although the first bounding volume may be a relatively accurate approximation of the actual
30 scene (e.g. room) being defined, the first bounding volume may not match exactly to the surface of the scene (e.g. room) being defined (e.g. where the, e.g. walls of the room may be uneven or have surface roughness). To allow for this when generating the texture, the points of the surface that are sampled for the texture (for each texel), can be, and are in an embodiment, allowed to, fall on, outside of or
35 inside the first bounding volume (rather than the sampling points being constrained

to be on the walls of the bounding volume). This will avoid holes being introduced within the texture where the first bounding volume does not match exactly with the scene geometry. The surface is in an embodiment sampled at a position as close to the corresponding position on the first bounding volume (in the relevant direction from the reference position for the texture) as possible, as the further a sampling point is from the first bounding volume, the greater the error that will be introduced when using the texture, particularly when making a local correction (as described above). Accordingly, the first bounding volume is in an embodiment defined such that it matches the scene (the surface defining the scene) closely (in an embodiment as closely as possible), such that the texture can be generated using samples near to (as near as possible to) the walls of the first bounding volume.

In an embodiment, as well as generating the texture, the first bounding volume that the texture represents (is to be used with) is also generated and stored. In an embodiment data defining the first bounding volume is generated and stored in association with (associated with) the texture. Again, this bounding volume information can be subject to any desired post-processing operations, such as compression, if desired.

In an embodiment, a graphics texture is generated and stored, in an embodiment together with an associated "first" bounding volume", e.g. for each scene that is anticipated may be displayed when executing the application, e.g. game, that the texture and scene relate to. Correspondingly, an array of pre-determined intersection position-indicating information (texture) is in an embodiment generated and stored, in an embodiment together with an associated "second" bounding volume, e.g. for each scene. In an embodiment, plural graphics textures are generated and stored, e.g. for respective scenes that is anticipated may be displayed when executing the application, e.g. game, that the textures and scenes relate to. Correspondingly, plural arrays of pre-determined intersection position-indicating information (textures) are in an embodiment generated and stored, e.g. for the respective scenes.

For example, where a game will include plural scenes, such as rooms, in an embodiment respective pairs of pre-determined intersection position-indicating textures and property-indicating textures (cubemaps) are generated for each scene (e.g. room) that could fall to be displayed when the game is played. The textures can then be stored, e.g., with the remaining game data for use when the game is being executed.

Any or all of the above processes could be performed for sampling points individually, or some or all of them could be performed for sets of plural sampling points (which would then have the effect of performing the process for the sampling points within the set that is being considered). For example, where the graphics processing unit operates on fragments that each represent a set of plural sampling positions, the process of the technology described herein could be performed on a per-fragment basis, rather than individually for each sampling point that a fragment represents (and in an embodiment this is what is done). In this case, there would accordingly be, e.g., a single property determined for a fragment, which would then be used for each sampling point that the fragment is being used to render.

It will also be appreciated that although embodiments have been described above with particular reference to a given sampling point, the technique of the technology described herein can be, and in an embodiment is, used for plural sampling points, and in an embodiment for each sampling point that needs to be considered when rendering the scene.

Thus, the process is in an embodiment repeated for each sampling point for a primitive that is on or within the (first) bounding volume that is being considered, and for each primitive that is on or within the (first) bounding volume that is being considered.

Similarly, the process is in an embodiment repeated for plural scenes in a sequence of scenes being rendered, e.g., and in an embodiment, for each frame of a sequence of frames to be rendered.

The technology described herein can be used in and with any suitable and desired graphics processing system and unit. The graphics processing unit in an embodiment comprises a graphic processing pipeline.

The technology described herein may be used with immediate mode renderers or tiled renderers (tile-based graphics processing systems). Thus, in one embodiment, the graphics processing unit is a tiled-based graphics processing unit.

The graphic processing unit in an embodiment comprise a plurality of processing stages, including at least a rasteriser that rasterises input primitives to generate graphics fragments to be processed, each graphics fragment having one or more sampling positions associated with it, and a renderer that processes fragments generated by the rasteriser to generate output fragment data.

The rasteriser of the graphics processing unit will in an embodiment generate graphics fragments to be rendered to generate rendered graphics data for

sampling points of the desired graphics output, such as a frame to be displayed. Each graphics fragment that is generated by the rasteriser has associated with it a set of sampling points of the graphics output and is to be used to generate rendered graphics data for one or more of the sampling points of the set of sampling points associated with the fragment.

The rasteriser may be configured to generate the fragments for rendering in any desired and suitable manner. It will in an embodiment receive e.g. primitives to be rasterised, test those primitives against sets of sampling point positions, and generate fragments representing the primitives accordingly.

The renderer should process the fragments generated by the rasteriser to generate rendered fragment data for (covered) sampling points that the fragments represent. These rendering processes may include, for example, fragment shading, blending, texture-mapping, etc.. In an embodiment the renderer is in the form of or includes a programmable fragment shader.

The graphics processing unit may also contain any other suitable and desired processing stages that a graphics processing unit may contain such as an early depth (or an early depth and stencil) tester, a late depth (or depth and stencil) tester, a blender, a tile buffer or buffers, a write-out unit, etc..

The graphics processing unit in an embodiment also comprises, and/or is in communication with, one or more memories and/or memory devices that store the data described herein, such as the textures, scene bounding volumes, etc., and/or that store software for performing the processes described herein. The graphics processing unit may also be in communication with the host microprocessor, and/or with a display for displaying images based on the data generated by the graphics processing unit.

The graphics processing unit in an embodiment includes at least local memory, such as (on-chip) buffer or buffers, and/or register(s), that can be used to store the data required for the intersection position determination process and the determined intersection position. Where present, buffers can be used for this purpose, if desired. The graphics processing unit can in an embodiment also cache determined intersection position values for future use, if desired.

The technology described herein can be used for all forms of output that a graphics processing unit may be used to generate, such as frames for display, render-to-texture outputs, etc..

In an embodiment, the various functions of the technology described herein are carried out on a single graphics processing platform that generates and outputs the rendered fragment data that is, e.g., written to the frame buffer for the display device.

5 The technology described herein can be implemented in any suitable system, such as a suitably configured micro-processor based system. In an embodiment, the technology described herein is implemented in a computer and/or micro-processor based system.

10 The various functions of the technology described herein can be carried out in any desired and suitable manner. For example, the functions of the technology described herein can be implemented in hardware or software, as desired. Thus, for example, unless otherwise indicated, the various functional elements and "means" of the technology described herein may comprise a suitable processor or processors, controller or controllers, functional units, circuitry, processing logic, 15 microprocessor arrangements, etc., that are operable to perform the various functions, etc., such as appropriately dedicated hardware elements and/or programmable hardware elements that can be programmed to operate in the desired manner.

20 It should also be noted here that, as will be appreciated by those skilled in the art, the various functions, etc., of the technology described herein may be duplicated and/or carried out in parallel on a given processor. Equally, the various processing stages may share processing circuitry, etc., if desired.

25 Furthermore, any one or more or all of the processing stages of the technology described herein may be embodied as processing stage circuitry, e.g., in the form of one or more fixed-function units (hardware) (processing circuitry), and/or in the form of programmable processing circuitry that can be programmed to perform the desired operation. Equally, any one or more of the processing stages and processing stage circuitry of the technology described herein may be provided as a separate circuit element to any one or more of the other processing stages or 30 processing stage circuitry, and/or any one or more or all of the processing stages and processing stage circuitry may be at least partially formed of shared processing circuitry.

Subject to any hardware necessary to carry out the specific functions discussed above, the graphics processing system can otherwise include any one or

more or all of the usual functional units, etc., that graphics processing systems include.

It will also be appreciated by those skilled in the art that all of the described embodiments of the technology described herein can, and in an embodiment do, include, as appropriate, any one or more or all of the optional features described herein.

The methods in accordance with the technology described herein may be implemented at least partially using software e.g. computer programs. It will thus be seen that when viewed from further embodiments the technology described herein comprises computer software specifically adapted to carry out the methods herein described when installed on a data processor, a computer program element comprising computer software code portions for performing the methods herein described when the program element is run on a data processor, and a computer program comprising code adapted to perform all the steps of a method or of the methods herein described when the program is run on a data processing system. The data processor may be a microprocessor system, a programmable FPGA (field programmable gate array), etc..

The technology described herein also extends to a computer software carrier comprising such software which when used to operate a graphics processor, renderer or microprocessor system comprising a data processor causes in conjunction with said data processor, said processor, renderer or system to carry out the steps of the methods of the technology described herein. Such a computer software carrier could be a physical storage medium such as a ROM chip, CD ROM, RAM, flash memory, or disk, or could be a signal such as an electronic signal over wires, an optical signal or a radio signal such as to a satellite or the like.

It will further be appreciated that not all steps of the methods of the technology described herein need be carried out by computer software and thus from a further broad embodiment the technology described herein comprises computer software and such software installed on a computer software carrier for carrying out at least one of the steps of the methods set out herein.

The technology described herein may accordingly suitably be embodied as a computer program product for use with a computer system. Such an implementation may comprise a series of computer readable instructions either fixed on a tangible, non-transitory medium, such as a computer readable medium, for example, diskette, CD-ROM, ROM, RAM, flash memory, or hard disk. It could

also comprise a series of computer readable instructions transmittable to a computer system, via a modem or other interface device, over either a tangible medium, including but not limited to optical or analogue communications lines, or intangibly using wireless techniques, including but not limited to microwave, infrared or other transmission techniques. The series of computer readable instructions embodies all or part of the functionality previously described herein.

Those skilled in the art will appreciate that such computer readable instructions can be written in a number of programming languages for use with many computer architectures or operating systems. Further, such instructions may be stored using any memory technology, present or future, including but not limited to, semiconductor, magnetic, or optical, or transmitted using any communications technology, present or future, including but not limited to optical, infrared, or microwave. It is contemplated that such a computer program product may be distributed as a removable medium with accompanying printed or electronic documentation, for example, shrink-wrapped software, pre-loaded with a computer system, for example, on a system ROM or fixed disk, or distributed from a server or electronic bulletin board over a network, for example, the Internet or World Wide Web.

Embodiments of the technology described herein will now be described in the context of the processing of computer graphics for display.

When a computer graphics image is to be displayed, it is usually first defined as a series of primitives (polygons), which primitives are then divided (rasterised) into graphics fragments for graphics rendering in turn. During a normal graphics rendering operation, the renderer will modify the (e.g.) colour (red, green and blue, RGB) and transparency (alpha, a) data associated with each fragment so that the fragments can be displayed correctly. Once the fragments have fully traversed the renderer, then their associated data values are stored in memory, ready for output for display.

Figure 1 shows schematically a graphics processing pipeline 1 that may operate in accordance with the technology described herein. The graphics processing pipeline 1 is a tiled deferred renderer with a fully programmable GPGPU context, for example a renderer which executes partly via Direct Compute, OpenCL, CUDA, etc..

As the graphics processing pipeline 1 shown in Figure 1 is a tile-based renderer, it will thus produce tiles of a render output data array, such as an output frame to be generated.

(In tile-based rendering, rather than the entire render output, e.g., frame, effectively being processed in one go as in immediate mode rendering, the render output, e.g., frame to be displayed, is divided into a plurality of smaller sub-regions, usually referred to as "tiles". Each tile (sub-region) is rendered separately (typically one-after-another), and the rendered tiles (sub-regions) are then recombined to provide the complete render output, e.g., frame for display. In such arrangements, the render output is typically divided into regularly-sized and shaped sub-regions (tiles) (which are usually, e.g., squares or rectangles), but this is not essential. Each tile corresponds to a respective set of screen space sampling positions.)

The render output data array may typically be an output frame intended for display on a display device, such as a screen or printer, but may also, for example, comprise intermediate data intended for use in later rendering passes (also known as a "render to texture" output), etc..

Figure 1 shows the main elements and pipeline stages of the graphics processing pipeline 1 that are relevant to the operation of the present embodiment. As will be appreciated by those skilled in the art there may be other elements of the graphics processing pipeline that are not illustrated in Figure 1. It should also be noted here that Figure 1 is only schematic, and that, for example, in practice the shown functional units and pipeline stages may share significant hardware circuits, even though they are shown schematically as separate stages in Figure 1. It will also be appreciated that each of the stages, elements and units, etc., of the graphics processing pipeline as shown in Figure 1 may be implemented as desired and will accordingly comprise, e.g., appropriate circuitry and/or processing logic, etc., for performing the necessary operation and functions.

Figure 1 shows schematically the pipeline stages after the graphics primitives (polygons) 2 for input to the rasterisation process have been generated. Thus, at this point the graphics data (the vertex data) has undergone fragment frontend operations 8, such as transformation operations (not shown), vertex shading, interpolation and a primitive set-up stage (not shown) to set-up the primitives to be rendered, in response to the commands and vertex data provided to the graphics processor.

- 27 -

As shown in Figure 1, this part of the graphics processing pipeline 1 includes a number of stages, including a rasterisation stage 3, an early Z (depth) and stencil test stage 4, a renderer in the form of a fragment shading stage 6, a late Z (depth) and stencil test stage 7, a blending stage 9, a tile buffer 10 and a
5 downsampling and writeout (multisample resolve) stage 13.

The rasterisation stage 3 of the graphics processing pipeline 1 operates to rasterise the primitives making up the render output (e.g. the image to be displayed) into individual graphics fragments for processing. To do this, the rasteriser 3 receives graphics primitives 2 for rendering, rasterises the primitives to sampling
10 points and generates graphics fragments having appropriate positions (representing appropriate sampling positions) for rendering the primitives.

The fragments generated by the rasteriser are then sent onwards to the rest of the pipeline for processing.

The early Z/stencil stage 4 performs a Z (depth) test on fragments it
15 receives from the rasteriser 3, to see if any fragments can be discarded (culled) at this stage. To do this, it compares the depth values of (associated with) fragments issuing from the rasteriser 3 with the depth values of fragments that have already been rendered (these depth values are stored in a depth (Z) buffer that is part of the tile buffer 10) to determine whether the new fragments will be occluded by
20 fragments that have already been rendered (or not). At the same time, an early stencil test is carried out.

Fragments that pass the fragment early Z and stencil test stage 4 are then sent to the fragment shading stage 6. The fragment shading stage 6 performs the appropriate fragment processing operations on the fragments that pass the early Z
25 and stencil tests, so as to process the fragments to generate the appropriate rendered fragment data.

This fragment processing may include any suitable and desired fragment shading processes, such as executing fragment shader programs on the fragments, applying textures to the fragments, applying fogging or other operations to the
30 fragments, etc., to generate the appropriate fragment data. In the present embodiment, the fragment shading stage 6 is in the form of a shader pipeline (a programmable fragment shader), but other arrangements, such as the use also or instead of fixed function fragment shading units would be possible, if desired.

There is then a "late" fragment Z and stencil test stage 7, which carries out,
35 *inter alia*, an end of pipeline depth test on the shaded fragments to determine

whether a rendered fragment will actually be seen in the final image. This depth test uses the Z-buffer value for the fragment's position stored in the Z-buffer in the tile buffers 10 to determine whether the fragment data for the new fragments should replace the fragment data of the fragments that have already been rendered, by
5 comparing the depth values of (associated with) fragments issuing from the fragment shading stage 6 with the depth values of fragments that have already been rendered (as stored in the depth buffer). This late fragment depth and stencil test stage 7 also carries out any necessary "late" alpha and/or stencil tests on the fragments.

10 The fragments that pass the late fragment test stage 7 are then subjected to, if required, any necessary blending operations with fragments already stored in the tile buffer 10 in the blender 9. Any other remaining operations necessary on the fragments, such as dither, etc. (not shown) are also carried out at this stage.

Finally, the (blended) output fragment data (values) are written to the tile
15 buffer 10 from where they can, for example, be output to a frame buffer for display. The depth value for an output fragment is also written appropriately to a Z-buffer within the tile buffer 10. (The tile buffer stores colour and depth buffers that store appropriate colour, etc., values or a Z-value, respectively, for each sampling position that the buffers represent (in essence for each sampling position of a tile that is being processed).) These buffers store an array of fragment data that
20 represents part (a tile) of the overall render output (e.g. image to be displayed), with respective sets of sample values in the buffers corresponding to respective pixels of the overall render output (e.g. each 2x2 set of sample values may correspond to an output pixel, where 4x multisampling is being used).

25 The tile buffer is provided as part of RAM that is located on (local to) the graphics processing pipeline (chip).

The data from the tile buffer 10 is input to a downsampling (multisample resolve) write out unit 13, and thence output (written back) to an external memory output buffer, such as a frame buffer of a display device (not shown). (The display
30 device could comprise, e.g., a display comprising an array of pixels, such as a computer monitor or a printer.)

In the present embodiments, the downsampling and writeout unit 13 downsamples (in either a fixed or variable fashion) the fragment data stored in the tile buffer 10 to the appropriate resolution for the output buffer (device) (i.e. such

that an array of pixel data corresponding to the pixels of the output device is generated), to generate output values (pixels) for output to the output buffer.

Once a tile of the render output has been processed and its data exported to a main memory (e.g. to a frame buffer in a main memory (not shown)) for storage, the next tile is then processed, and so on, until sufficient tiles have been processed to generate the entire render output (e.g. frame (image) to be displayed). The process is then repeated for the next render output (e.g. frame) and so on.

Other arrangements for the graphics processing pipeline 1 would, of course, be possible.

The above describes certain features of the operation of the graphics processing system shown in Figure 1. Further features of the operation of the graphics processing system shown in Figure 1 that allow the effect of reflections in an image being rendered to be simulated in accordance with embodiments of the technology described herein will now be described.

The present embodiments operate to simulate the effect of reflections within a defined volume for a scene. Figure 2 illustrates this and shows an exemplary “first” bounding volume 20 in the form of a bounding box that is defined in world space, that represents a volume within the scene being rendered. In the present embodiment it is assumed that the bounding volume 20 corresponds to the entire scene and represents e.g., a room, but other arrangements would, of course, be possible. As shown in Figure 2, it is also assumed that a reflective surface 21 is present in the scene.

The first bounding volume 20 should be a convex bounding volume, i.e. such that for any two points within the bounding volume, every point on the straight line segment that joins the two points is also within the first bounding volume. However, it would also be possible for the first bounding volume to comprise a shape that is not “strictly” convex in this sense. For example, the first bounding volume could comprise a substantially convex volume having one or more holes 29, which could be used to represent windows in the environment (e.g. room) of the scene being rendered.

In the present embodiment, a “second” bounding volume 50 that bounds the first bounding volume 20 is defined in world space. The second bounding volume 50 is effectively a simplified version of the first bounding volume 20. In Figure 2, the second bounding volume 50 takes the form of a bounding sphere, but other arrangements, such as a bounding box, would be possible. In general, any volume

where the intersection point of a vector on the volume can be analytically calculated in a relatively simple manner may be used.

A distance-indicating texture is generated that stores distance values representing the distance from the surface of the first bounding volume 20 to the surface of the second bounding volume 50 with respect to various points on the second bounding volume 50 and in various directions. In the present embodiment, each distance value in the texture comprises the distance D from the position on the second bounding volume 50 intersected by a vector to the position on the first bounding volume 20 intersected by the vector for a given combination of the position on the second bounding volume 50 intersected by the vector and the direction of the vector, although other arrangements would be possible.

In the present embodiment, each distance value in the distance texture is in the form of a floating point number, and has a certain precision such as a 16-bit precision, etc., although other arrangements would be possible.

As can be seen from Figure 2, where the scene to be rendered is a 3D scene, the position on the second bounding volume 50 intersected by a vector and the direction of the vector can each be defined by a 2D direction (e.g. a pair of spherical coordinates), and so the texture may comprise a 4D texture. Where the scene is a 2D scene, such that the position on the second bounding volume intersected by the vector and the direction of the vector can each be defined by a 1D direction (e.g. an angle), the texture may comprise a 2D texture.

It would also be possible, where the scene is a 3D scene, for the texture to comprise a 3D texture. In this case, a 4D texture having a resolution of, for example, 32x32x16x16 data elements, e.g. corresponding to 32x32 possible intersection positions on the second bounding volume and 16x16 possible vector directions, could be “flattened” into a 3D texture having a resolution of 32x32x256 data elements.

Figure 3 shows schematically the generation of the distance-indicating texture. In the present embodiment, the distance-indicating texture is generated by firstly defining the first bounding volume 20 as a “watertight” polygonal mesh (step 60), defining the second bounding volume 50 that bounds the first bounding volume 20 (step 61), defining the resolution of the texture (e.g. 32x32x16x16), and defining how the texture maps to the second bounding volume 50 (i.e. defining the reference position for the second bounding volume 50 and how each position on the second bounding volume corresponds to each spherical coordinate, etc.) (step 62). Then,

a 4D texture is created (step 63), and for each of the (32x32) possible intersection positions on the second bounding volume 50, for each of the (16x16) possible vector directions, a ray is intersected with the first bounding volume 20, and (where possible) a distance value is calculated. As the first bounding volume 20 is defined as a watertight polygonal mesh, this intersection may be done by iterating through each polygon in the mesh and performing a ray-polygon intersection calculation. Thus, as shown in Figure 3, the process involves iteratively determining, for each intersection position and for each direction, whether a ray intersects with the polygonal mesh (steps 64-69).

Each calculated distance value is then stored in a 4D texture, at the texel corresponding to the intersection position on the second bounding volume 50 and the vector direction in question (step 70). If, for a particular intersection position and vector direction, there is no intersection with the first bounding volume 20 (e.g. where the first and second bounding volumes are non-identical), a "special" value is stored in the texture that indicates there is no intersection (step 71). For a floating point texture, this could be, for example, positive infinity. It would also be possible, rather than store a special value, to store the distance value from the closest texel where an intersection does occur.

Once the distance texture has been generated, it may optionally be subjected to compression. Any suitable texture compression scheme can be used to compress the distance texture. Where the texture is a 3D texture, "standard" 3D texture compression schemes, such as the Adaptive Scalable Texture Compression (ASTC) scheme, can be used. Since in the present embodiment only an approximate intersection point is determined, relatively heavy compression may be used, reducing the storage and bandwidth requirements.

The distance texture generation and any compression of the distance texture can be performed as desired and using any suitable processing system. In an embodiment the distance texture is generated "offline" and stored, and then provided to the graphics processing pipeline for use when it is required.

In the present embodiment a cube texture corresponding to the surface of the bounding volume 20 that stores colour (RGB) values representing the colour of the surface of the bounding volume 20 is also generated. Figure 4 shows the corresponding cube texture 30 for the bounding volume 20 shown in Figure 2. The cube texture 30 will store colour values for the bounding volume surface.

Figure 5 shows schematically the generation of the cube texture shown in Figure 4. As shown in Figure 5, the process starts by rendering the scene to the cube texture from the reference position 31 (shown in Figure 2) that the cube texture is to be defined with respect to (step 40). (The cube texture 30 will be defined with respect to a reference position 31 (shown in Figure 2) within the volume that the cube texture represents (thus the bounding volume for the scene), and will be sampled by determining the direction from that reference position that the sampled value (the sampled position) is required for.)

This process uses as its inputs information about the scene 41, information defining the first bounding volume 20, information defining the reference position 31 for the cube texture, and information indicating the resolution size 42 that is desired for the cube texture.

The cube texture 30 is generated by rendering an image that will represent the surface of the first bounding volume 20 from the viewpoint of the reference position 31 for the texture. To do this respective positions on the surface (the image representing the surface) for the scene (as defined by the scene information 41) are sampled for respective positions on the first bounding volume (in the relevant direction from the reference position 31 for the texture). In this process, as the first bounding volume may be an approximation of the actual scene and so may not match exactly to the surface of the scene (e.g. room) being defined, the points of the surface that are sampled for the texture (for each texel) do not need to fall exactly on the walls of the bounding volume 20 but can be on, outside of or inside the bounding volume. The output 47 of this cube texture generation process is then the cube texture as shown in Figure 4, which represents the surface of the bounding volume for the scene that the cube texture relates to at respective positions on the surface of the bounding volume.

As shown in Figure 5, once the cube texture has been generated, it may optionally be subjected to various processing effects and optimisations, such as blurring, gamma correction, brightness and/or contrast enhancement, etc. (step 43). This step uses as inputs desired post-processing effect parameters, such as blurring parameters, gamma correction parameters, contrast parameters, etc..

Furthermore, as shown in Figure 5, the generated (and if desired post-processed) cube texture can be compressed (step 45), using desired compression parameters 46. Any suitable texture compression scheme can be used to compress the cube texture.

The cube texture generation and any post-processing/compression of the cube texture can be performed as desired and using any suitable processing system. In an embodiment the cube texture is generated “offline” and stored, and then provided to the graphics processing pipeline for use when it is required.

5 In an embodiment, as well as generating the textures, the first and second bounding volumes that the textures represent (are to be used with) are also generated and stored. In an embodiment data defining the first and second bounding volumes is generated and stored in association with (associated with) the textures. Again, this bounding volume information can be subject to any desired
10 post-processing operations, such as compression, if desired.

As discussed above, the present embodiment relates to arrangements for determining the effects of reflections, e.g. from the reflective surface 21 as seen from the viewpoint (camera) position 22 for the scene. The use of the distance texture and the cube texture shown in Figure 4 to determine the effect of reflections
15 within the bounding volume 20 will now be described with reference to Figure 2.

The present embodiment operates to determine for respective sampling points on the reflective surface 21, the effect of reflections using the cube texture 30. Figure 2 shows the process for a given sampling point 23. This process is repeated for each sampling point on the reflective surface 21.

20 The process starts by determining the reflected vector 24 for the sampling point 23 that is being considered. This process is based on the law of reflection, and uses as its input the vector 25 from the viewpoint (camera) position 22 for the scene to the sampling point 23, and the normal 26 of the surface at the sampling point 23.

25 The next step is to apply a “local correction” to the reflected vector 24 to ensure that the correct position within the cube texture 30 is sampled. This is required, because, as can be seen from Figure 2, for example, the reference position 31 that the cube texture is sampled from (with respect to) may not (and typically will not) correspond to the sampling position 23 being considered, such
30 that simply taking the reflected vector 24 and using that to sample from the reference position 31 for the cube texture 30 will not sample the correct part of the cube texture.

To apply this “local correction”, as shown in Figure 2, the intersection point 27 on the surface of the bounding volume 20 of the reflected vector 24 is firstly
35 (approximately) determined. As will be described in more detail below, this is done

- 34 -

by determining and using the intersection point 51 of the reflected vector 24 on the surface of a bounding sphere 50 that encompasses the first bounding volume 20. The vector 28 from the reference position 31 that the cube texture 30 is defined with respect to to the intersection point 27 on the first bounding volume 20 is
5 determined, and that vector 28 is then used to sample the cube texture 30. This will, accordingly, and as can be seen from Figure 2, sample the cube texture 30 at the intersection position 27 that corresponds to the portion of the surface of the bounding volume 20 that is intersected by the reflected vector 24.

10 The cube texture 30 is then sampled at the position corresponding to the intersection position 27 to determine the colour and/or transparency of the surface of the bounding volume 20 at that point. The cube texture 30 can be sampled using any desired texture sampling (filtering) process, such as bilinear and/or tri-linear filtering (e.g. in the case where the cube texture 30 is in the form of mipmaps). The filtering may, e.g., be provided by a texturing module (stage) of the graphics
15 processing pipeline. This then gives a colour value that can be used to simulate the effects of reflection at the sampling point 23.

This process is then be repeated for the sampling points for all the primitives on the reflective surface 21 that are to be rendered (and then for the next frame (if appropriate), and so on).

20 The use of the intersection point 51 of the reflected vector 24 on the surface of the bounding sphere 50 and the distance texture associated with the bounding sphere 50 to approximately determine the intersection point 27 on the surface of the first bounding volume 20 of the reflected vector 24 will now be described in more detail with reference to Figure 6. Figure 6 shows a schematic two-dimensional
25 representation of the scene, where the first bounding volume 20 is shown a triangle and the bounding sphere 50 is shown as a circle that encompasses the bounding volume 20.

In the present embodiment, the intersection point 51 on the surface of the bounding sphere 50 of the reflected vector 24 from the sampling point 23 is used to
30 determine the unit direction vector a between the centre point 52 of the bounding sphere 50 and the intersection point 51. This intersection point 51 is analytically calculated using as inputs the radius and the centre point 52 of the bounding sphere 50, and the origin 23 and direction of the vector 24. The unit direction vector b of the reflected vector 24 is also determined.

The unit direction vectors a and b are then converted into respective pairs of spherical coordinates, and the spherical coordinates are used to look up a distance value from the distance texture. The first two texture coordinates (x , y) may be the spherical coordinates of the direction of the vector a between the centre 52 and intersection point 51 of the bounding sphere 50, and the next two texture coordinates (z , w) may be the spherical coordinates of the direction b of the reflected vector 24. The looked up distance value will (approximately) correspond to the distance d from the position 51 on the second bounding volume 50 intersected by the vector 24 to the position 27 on the first bounding volume 20 intersected by the vector 24 for the particular combination of the position 51 on the second bounding volume 50 intersected by the vector 24 and the direction of the vector 24 in question.

Where the texture comprises a 2D or 4D texture, the texture look up may be a 2D or 4D texture look up. Where the texture comprises a 3D texture comprising flattened 4D distance information (e.g. where there is no hardware support for 4D textures) two 3D texture lookups may be performed and results may be appropriately interpolated.

Using the determined distance d , the approximate coordinates of the position 27 on the first bounding volume 20 intersected by the vector 24 can then be calculated by multiplying the unit direction vector of the vector 24 with the distance obtained from the texture, and subtracting this resultant vector from the position 51 on the second bounding volume 50 intersected by the vector 24.

Various modifications, additions and alternatives to the above-described embodiments of the technology described herein would be possible, if desired.

In particular, although the above embodiment has been described in terms of taking account of the effects of reflections in a scene being rendered, the techniques of the technology described herein can be used in a number of other arrangements.

For example, the techniques of the technology described herein may be used in other environment mapping techniques, such as the Applicant's shadow mapping techniques described in GB 1413145.2, and the Applicant's refraction mapping techniques described in GB 1503391.3 and GB 1503386.3.

The techniques of the technology described herein may also be used to detect collisions between particles and an environment, e.g. for in-game effects. In this case, the intersection position of a particle velocity vector on a bounding

volume may be determined and used to determine whether the particle will collide with the environment, e.g. within a given time period.

Although the above embodiment has been described in terms of the use of a distance-indicating texture that includes plural distance values to determine the position 27 on the first bounding volume 20 intersected by the vector 24, it would also be possible for the texture to instead include plural position values (e.g. coordinates) that (approximately) correspond to the coordinates of the position 27 on the first bounding volume 20 intersected by the vector 24 for the particular combination of the position 51 on the second bounding volume 50 intersected by the vector 24 and the direction of the vector 24 in question. In this case, the spherical coordinates can be used to look up the position value from the texture. As will be appreciated by those having skill in the art, this would then reduce the number of calculations required at run-time (since it would not then be necessary to calculate the coordinates of the position 27 using the distance d), but at the expense of requiring more data to be stored in the texture and increased memory bandwidth.

As can be seen from the above, the technology described herein, in its embodiments at least, provides an effective, and processing efficient, mechanism for determining the intersection point of a vector on a volume, that can allow higher quality environment mapping rendering to be achieved in a bandwidth and processing efficient manner.

This is achieved, in the embodiments of the technology described herein at least, by determining the position intersected by the vector on a second bounding volume associated with the first bounding volume, using the determined position to determine information indicative of a position intersected on the first bounding volume by the vector, and using the determined information to determine a position on the first bounding volume intersected by the vector.

The foregoing detailed description has been presented for the purposes of illustration and description. It is not intended to be exhaustive or to limit the technology to the precise form disclosed. Many modifications and variations are possible in the light of the above teaching. The described embodiments were chosen in order to best explain the principles of the technology and its practical application, to thereby enable others skilled in the art to best utilise the technology in various embodiments and with various modifications as are suited to the

- 37 -

particular use contemplated. It is intended that the scope be defined by the claims appended hereto.

CLAIMS

What is claimed is:

- 5 1. A method of operating a graphics processing system when rendering a scene for output, in which a first bounding volume representative of the volume of all or part of the scene to be rendered is defined; the method comprising:
- determining an intersection position for a vector with the first bounding volume by:
- 10 determining the position intersected by the vector on a second bounding volume associated with the first bounding volume;
- using the determined position to determine information indicative of an intersection position for the first bounding volume and the vector; and
- using the determined information to determine an intersection
- 15 position for the vector with the first bounding volume.
2. The method of claim 1, wherein the first bounding volume comprises a convex volume.
- 20 3. The method of claim 1 or 2, wherein the second bounding volume at least partially bounds and/or is at least partially bounded by the first bounding volume.
4. The method of claim 1, 2 or 3, wherein the second bounding volume has a form that is relatively simple compared to the form of the first bounding volume.
- 25 5. The method of any one of the preceding claims, wherein the direction of the vector is used in determining the information indicative of an intersection position for the first bounding volume and the vector.
- 30 6. The method of any one of the preceding claims, wherein the information indicative of an intersection position for the first bounding volume and the vector is determined by using the determined position to look up the information from pre-determined information.

7. The method of claim 6, wherein the pre-determined information comprises information indicative of an intersection position on the first bounding volume for a vector for each of plural combinations of the position on the second bounding volume intersected by the vector and the direction of the vector.

5

8. The method of claim 6 or 7, wherein the pre-determined information comprises a 3D array of 4D pre-determined information.

9. The method of claim 6, 7 or 8, wherein the pre-determined information comprises a graphics texture, and the information indicative of an intersection position for the first bounding volume and the vector is looked up by sampling the graphics texture.

10. The method of any one of the preceding claims, wherein the information indicative of an intersection position for the first bounding volume and the vector comprises position information.

11. The method of any one of claims 1-10, wherein the information indicative of an intersection position for the first bounding volume and the vector comprises distance information, and the method comprises:

using the determined position to determine a distance value indicative of the distance from the determined position intersected on the second bounding volume by the vector to an intersection position for the first bounding volume and the vector; and

using the determined distance value to determine an intersection position for the vector with the first bounding volume.

12. A method of generating a texture for use in a graphics processing system when rendering a scene for output in which a first bounding volume representative of the volume of all or part of the scene to be rendered is defined, the method comprising:

generating a graphics texture comprising an array of texture texels by:
setting the texel values in the texture such that each texel stores a value indicative of an intersection position on the first bounding volume for a vector for a particular combination of the position on a second bounding volume associated with

- 40 -

the first bounding volume intersected by the vector and the direction of the vector;
and

storing data representing the texel values of the texture.

5 13. A graphics processing unit comprising:

processing circuitry configured to determine an intersection position for a
vector with a first bounding volume representative of the volume of all or part of a
scene to be rendered by:

10 determining the position intersected by the vector on a second
bounding volume associated with the first bounding volume;
using the determined position to determine information indicative of
an intersection position for the first bounding volume and the vector; and
using the determined information to determine an intersection
position for the vector with the first bounding volume.

15

14. The graphics processing unit of claim 13, wherein the first bounding volume
comprises a convex volume.

20 15. The graphics processing unit of claim 13 or 14, wherein the second
bounding volume at least partially bounds and/or is at least partially bounded by the
first bounding volume.

25 16. The graphics processing unit of claim 13, 14 or 15, wherein the second
bounding volume has a form that is relatively simple compared to the form of the
first bounding volume.

30 17. The graphics processing unit of any one claims 13-16, wherein the direction
of the vector is used in determining the information indicative of an intersection
position for the first bounding volume and the vector.

35

18. The graphics processing unit of any one of claims 13-17, wherein the
information indicative of an intersection position for the first bounding volume and
the vector is determined by using the determined position to look up the information
from pre-determined information.

35

19. The graphics processing unit of claim 18, wherein the pre-determined information comprises information indicative of an intersection position on the first bounding volume for a vector for each of plural combinations of the position on the second bounding volume intersected by the vector and the direction of the vector.

5

20. The graphics processing unit of claim 18 or 19, wherein the pre-determined information comprises a 3D array of 4D pre-determined information.

21. The graphics processing unit of claim 18, 19 or 20, wherein the pre-determined information comprises a graphics texture, and the information indicative of an intersection position for the first bounding volume and the vector is looked up by sampling the graphics texture

10

22. The graphics processing unit of any one claims 13-21, wherein the information indicative of an intersection position for the first bounding volume and the vector comprises position information.

15

23. The graphics processing unit of any one claims 13-21, wherein the information indicative of an intersection position for the first bounding volume and the vector comprises distance information, and the processing circuitry is configured to:

20

use the determined position to determine a distance value indicative of the distance from the determined position intersected on the second bounding volume by the vector to an intersection position for the first bounding volume and the vector; and

25

use the determined distance value to determine an intersection position for the vector with the first bounding volume.

24. An apparatus for generating a texture for use in a graphics processing system when rendering a scene for output, in which a first bounding volume representative of the volume of all or part of the scene to be rendered is defined, the apparatus comprising processing circuitry configured to:

30

generate a graphics texture comprising an array of texture texels by:

setting the texel values in the texture such that each texel stores a value

indicative of an intersection position on the first bounding volume for a vector for a

35

- 42 -

particular combination of the position on a second bounding volume associated with the first bounding volume intersected by the vector and the direction of the vector; and

storing data representing the texel values of the texture.

5

25. A texture for use in a graphics processing system when rendering a scene for output, in which a first bounding volume representative of the volume of all or part of the scene to be rendered is defined, the texture comprising:

an array of texture texels, in which:

10

the texel values are set such that each texel stores a value indicative of an intersection position on the first bounding volume for a vector for a particular combination of the position on a second bounding volume associated with the first bounding volume intersected by the vector and the direction of the vector.

15

26. A computer program comprising software code adapted to perform the method of any one of claims 1 to 12 when the program is run on a data processing system.

20

27. A method of operating a graphics processing system substantially as herein described with reference to any one of the accompanying drawings.

25

28. A method of generating a texture for use in a graphics processing system substantially as herein described with reference to any one of the accompanying drawings.

29. A graphics processing unit substantially as herein described with reference to any one of the accompanying drawings.

30

30. An apparatus for generating a texture for use in a graphics processing system substantially as herein described with reference to any one of the accompanying drawings.

1/6

Fig. 1





Fig. 2



3/6

Fig. 3




Fig. 4




Fig. 5





Fig. 6



INTERNATIONAL SEARCH REPORT

International application No
PCT/GB2016/052432

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06T11/00 G06T15/04 G06T15/50
ADD.

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)
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 practicable, search terms used)

EPO-Internal, WPI Data, COMPENDEX, EMBASE, INSPEC

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	LÁSZLÓ SZIRMAY-KALOS ET AL: "Approximate Ray-Tracing on the GPU with Distance Impostors", COMPUTER GRAPHICS FORUM, vol. 24, no. 3, 1 September 2005 (2005-09-01), pages 695-704, XP055313808, GB ISSN: 0167-7055, DOI: 10.1111/j.1467-8659.2005.0m894.x section 1 "introduction"; page 695, left-hand column, paragraph 2 - page 696, left-hand column, paragraph 1 section 1 "introduction"; page 696, left-hand column, line 11, paragraph 2 - last line, last paragraph section 2 "localization of the environment map"; page 696, left-hand column, line 1, -/--	1-30



Further documents are listed in the continuation of Box C.



See patent family 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 application or patent 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

28 October 2016

Date of mailing of the international search report

07/11/2016

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Gauthier, J

INTERNATIONAL SEARCH REPORT

International application No
PCT/GB2016/052432

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	paragraph 1 - right-hand column, last line, last paragraph sections 2.1 and 2.2; page 696, right-hand column - page 698, left-hand column, paragraph 1 section 3 "environment mapping with distance impostors"; page 698, left-hand column, paragraph 1 - page 700, left-hand column, paragraph 2 section 6 "the complete rendering algorithm"; page 702, right-hand column, paragraph 1 - paragraph 3 figures 1-4,8,15 ----- Roberto Lopez Mendez: "Reflections based on local cubemaps in unity", ARM Mali Graphics 2014, 7 August 2014 (2014-08-07), pages 1-19, XP055313837, ARM Community Blog Retrieved from the Internet: URL:https://community.arm.com/groups/arm-mali-graphics/blog/2014/08/07/reflections-based-on-local-cubemaps [retrieved on 2016-10-25] page 2 - page 4 page 6 - page 8 page 10 - page 12 page 15 figures 1-3,5-7 ----- -/--	1-3,5-7, 9,10, 12-15, 17-19, 21,22, 24-30

INTERNATIONAL SEARCH REPORT

International application No

PCT/GB2016/052432

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	LÁSZLÓ SZIRMAY-KALOS ET AL: "Specular Effects on the GPU: State of the Art", COMPUTER GRAPHICS FORUM, vol. 28, no. 6, 1 September 2009 (2009-09-01), pages 1586-1617, XP055313840, GB ISSN: 0167-7055, DOI: 10.1111/j.1467-8659.2009.01350.x section 3 "environment mapping"; page 1593, right-hand column, line 1 - page 1595, right-hand column, line 2 section 4 "Ray tracing in the sampled geometry"; page 1595, right-hand column, paragraph 1 - page 1596, left-hand column, paragraph 1 section 4.1 "Ray tracing distance maps"; page 1596, left-hand column, paragraph 1 - page 1597, right-hand column, last paragraph section 4.2 "Single reflections or refractions using the sampled geometry"; page 1597, right-hand column, paragraph 1 - page 1598, right-hand column, last paragraph page 1599, right-hand column, last paragraph figures 13,15-19,23 -----	1-30
X	Sebastien Lagarde ET AL: "Local Image-based Lighting with Parallax-corrected Cubemap", GamesConnection, 30 November 2012 (2012-11-30), XP055235164, Retrieved from the Internet: URL:https://seblagarde.files.wordpress.com/2012/08/parallax_corrected_cubemap-gameconnection2012.pdf [retrieved on 2015-12-09] page 5 - page 6 page 9 - page 11 page 17 page 21 - page 25 page 28 - page 43 ----- -/--	1-3,5-7, 9,10, 12-15, 17-19, 21,22, 24-30

INTERNATIONAL SEARCH REPORT

International application No

PCT/GB2016/052432

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	Shengying Li, et al.: "Real-time Reflection using Ray Tracing with Geometry Field", Eurographics 2006 Conference, Short Papers, 04-08 september 2006, Vienna, Austria, 4 September 2006 (2006-09-04), XP055313860, Retrieved from the Internet: URL:https://www.researchgate.net/profile/Xianfeng_Gu/publication/247344252_Real-time_Reflection_using_Ray_Tracing_with_Geometry_Field/links/0a85e53921f0f93327000000.pdf [retrieved on 2016-10-25] section 1; page 29, left-hand column, paragraphs 3,4,6, 7 - right-hand column section 1; page 30, left-hand column, paragraph 2 section 2; page 30, left-hand column, paragraphs 1,3 - right-hand column sections 3.1 and 3.2; page 30, right-hand column - page 31, left-hand column section 5; page 31, paragraph 2 -----	1,12,13, 24,25, 27-30
X	US 2005/093857 A1 (WANG LIFENG [CN] ET AL) 5 May 2005 (2005-05-05) -----	12,24, 25,28,30
A	paragraph [0029] - paragraph [0035] paragraphs [0042], [0045] claims 1,2,8-11 figures 3,5,6 -----	6-9,11, 18-21, 23,27,29
A	US 2008/024491 A1 (SATHE RAHUL P [US] ET AL) 31 January 2008 (2008-01-31) figures 1,4,5,6,7 paragraphs [0012], [0013], [0017] paragraph [0019] - paragraph [0027] paragraph [0032] - paragraph [0036] paragraph [0051] -----	6-9,11, 12, 18-21, 23-25, 28,30

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/GB2016/052432

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2005093857 A1	05-05-2005	US 2005093857 A1	05-05-2005
		US 2007297646 A1	27-12-2007

US 2008024491 A1	31-01-2008	CN 101496067 A	29-07-2009
		EP 2047430 A1	15-04-2009
		JP 4842379 B2	21-12-2011
		JP 2009545075 A	17-12-2009
		US 2008024491 A1	31-01-2008
		WO 2008014127 A1	31-01-2008
