

(51) International Patent Classification:
G06T 15/60 (2006.01)

(21) International Application Number:

PCT/US2010/035676

(22) International Filing Date:

20 May 2010 (20.05.2010)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

61/180,389	21 May 2009 (21.05.2009)	US
12/718,866	5 March 2010 (05.03.2010)	US
61/312,656	10 March 2010 (10.03.2010)	US

(63) Related by continuation (CON) or continuation-in-part (CIP) to earlier application:

US	12/718,866 (CIP)
Filed on	5 March 2010 (05.03.2010)

(71) Applicant (for all designated States except US): **SONY COMPUTER ENTERTAINMENT AMERICA INC.**
[US/US]; 919 E. Hillsdale Boulevard, 2nd Floor, Foster City, California 94404 (US).

(72) Inventor; and

(75) Inventor/Applicant (for US only): **DIAMAND, Ben**
[US/US]; c/o Sony Computer Entertainment America Inc., 919 E. Hillsdale Boulevard, 2nd Floor, Foster City, California 94404 (US).(74) Agents: **WAWRZYNIAK, Richard, E.** et al.; Fitch, Even, Tabin & Flannery, 120 South La Salle Street, Suite 1600, Chicago, Illinois 60603-3406 (US).

(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, 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, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, 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, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, 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, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

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

[Continued on next page]

(54) Title: METHOD AND APPARATUS FOR RENDERING SHADOWS




FIG. 4

(57) Abstract: A method for rendering shadows including identifying shadow casters (120) and shadow receivers (130) in a scene (100). In one embodiment constructing a first 3D grid (600), projecting the shadow casters (120, 610) and growing cells of the first 3D grid including a projected caster (610). Constructing a second 3D grid (700) and growing each cell of the second 3D grid including at least a portion of one shadow receiver and for which a corresponding cell of the first 3D grid is grown. Forming a third 3D (800) grid by intersecting the first 3D grid with the second 3D grid and marking for rendering each shadow caster (610) that intersects the third 3D grid. In another embodiment, for visible shadow receivers establishing depth information in a small white buffer (1310), rendering the shadow casters into a depth map, blurring the small white buffer and upsampling the small white buffer into a full sized white buffer (1320) using depth information as a metric.



— *before the expiration of the time limit for amending the
claims and to be republished in the event of receipt of*

amendments (Rule 48.2(h))

METHOD AND APPARATUS FOR RENDERING SHADOWS

CROSS-REFERENCE TO RELATED APPLICATIONS

5 This application claims the benefit of U.S. Provisional Application No. 61/312,656, filed March 10, 2010, entitled “METHOD AND APPARATUS FOR RENDERING SHADOWS”, the entire content and disclosure of which is hereby incorporated by reference herein in its entirety.

 This application also claims the benefit of U.S. Provisional Application No. 61/180,389, filed May 21, 2009, entitled “METHOD AND APPARATUS FOR
10 RENDERING IMAGE BASED PROJECTED SHADOWS WITH MULTIPLE DEPTH AWARE BLURS”, the entire content and disclosure of which is hereby incorporated by reference herein in its entirety.

 This application is a continuation-in-part of U.S. Application No. 12/718,866
15 filed March 5, 2010, entitled “METHOD AND APPARATUS FOR RENDERING IMAGE BASED PROJECTED SHADOWS WITH MULTIPLE DEPTH AWARE BLURS”, now U.S. Patent No _____, the entire content and disclosure of which is hereby incorporated by reference herein in its entirety.

20 **BACKGROUND OF THE INVENTION**

1. Field of the Invention

 The present invention relates generally to computer graphics, and more specifically to rendering shadows in computer graphics.

25 2. Discussion of the Related Art

 Shadow mapping is a process by which shadows are added to three-dimensional (3D) computer graphics.

SUMMARY OF THE INVENTION

30 Several embodiments of the invention advantageously provide a computer

readable storage medium storing one or more computer programs adapted to cause a processor based system to execute steps comprising identifying a set of shadow casters and a set of shadow receivers in a frame of a scene that is lighted by at least a first light source, constructing a first three-dimensional (3D) grid, from the first light source's point of view, that fully encloses a first volume of the scene, comprising one or more shadow casters of the set of shadow casters that fall within the first volume of the scene, projecting each of the one or more shadow casters to generate one or more projected shadow casters, growing each cell of the first 3D grid that includes at least a portion of the one or more projected casters, constructing a second three-dimensional (3D) grid, from the first light source's point of view, that fully encloses the first volume of the scene, comprising one or more shadow receivers of the set of shadow receivers that fall within the first volume of the scene, growing each cell of the second 3D grid that includes at least a portion of at least one of the one or more shadow receivers and for which a corresponding cell of the first 3D grid has been grown, marking for rendering each shadow receiver for which a cell of the second 3D grid has been grown, forming a third 3D grid by intersecting the first 3D grid with the second 3D grid such that the third 3D grid comprises the grown cells of the first grid and the grown cells of the second grid that intersect and marking for rendering into a shadow map each shadow caster that intersects the third 3D grid.

Another embodiment of the present invention provides a method comprising identifying a set of shadow casters and a set of shadow receivers in a frame of a scene that is lighted by at least a first light source, constructing a first three-dimensional (3D) grid, from the first light source's point of view, that fully encloses a first volume of the scene, comprising one or more shadow casters of the set of shadow casters that fall within the first volume of the scene, projecting each of the one or more shadow casters to generate one or more projected shadow casters, growing each cell of the first 3D grid that includes at least a portion of the one or more projected casters, constructing a second three-dimensional (3D) grid, from the first light source's point of view, that fully encloses the first volume of the scene, comprising one or more shadow receivers of the set of shadow receivers that fall within the first volume of the scene, growing each cell of the second 3D grid that includes at least a portion of at least one of the one or more shadow receivers and for which a corresponding cell of

the first 3D grid has been grown, marking for rendering each shadow receiver for which a cell of the second 3D grid has been grown, forming a third 3D grid by intersecting the first 3D grid with the second 3D grid such that the third 3D grid comprises the grown cells of the first grid and the grown cells of the second grid that intersect and marking for rendering into a shadow map each shadow caster that intersects the third 3D grid.

Yet another embodiment provides an apparatus, comprising a display and a processor based system coupled to the display, wherein the processor based system is configured to execute steps comprising identifying a set of shadow casters and a set of shadow receivers in a frame of a scene that is lighted by at least a first light source, constructing a first three-dimensional (3D) grid, from the first light source's point of view, that fully encloses a first volume of the scene, comprising one or more shadow casters of the set of shadow casters that fall within the first volume of the scene, projecting each of the one or more shadow casters to generate one or more projected shadow casters, growing each cell of the first 3D grid that includes at least a portion of the one or more projected casters, constructing a second three-dimensional (3D) grid, from the first light source's point of view, that fully encloses the first volume of the scene, comprising one or more shadow receivers of the set of shadow receivers that fall within the first volume of the scene, growing each cell of the second 3D grid that includes at least a portion of at least one of the one or more shadow receivers and for which a corresponding cell of the first 3D grid has been grown, marking for rendering each shadow receiver for which a cell of the second 3D grid has been grown, forming a third 3D grid by intersecting the first 3D grid with the second 3D grid such that the third 3D grid comprises the grown cells of the first grid and the grown cells of the second grid that intersect and marking for rendering into a shadow map each shadow caster that intersects the third 3D grid.

Another embodiment provides a computer readable storage medium storing one or more computer programs adapted to cause a processor based system to execute steps comprising identifying one or more groups of objects in a frame of content, wherein each group of objects includes one or more shadow casters and corresponding shadow receivers, for visible shadow receivers in the identified groups of objects, establishing depth information in a small white buffer, for each of the

identified groups of objects, rendering the shadow casters into a single depth map, blurring the small white buffer and upsampling the small white buffer into a full sized white buffer, wherein the upsampling comprises using depth information as a metric.

Yet another embodiment provides a method, comprising identifying one or
5 more groups of objects in a frame of content, wherein each group of objects includes one or more shadow casters and corresponding shadow receivers, for visible shadow receivers in the identified groups of objects, establishing depth information in a small white buffer, for each of the identified groups of objects, rendering the shadow casters
10 into a single depth map, blurring the small white buffer, upsampling the small white buffer into a full sized white buffer, wherein the upsampling comprises using depth information as a metric and is performed by a processor-based apparatus.

Further yet, another embodiment provides an apparatus, comprising a display and a processor based system coupled to the display, wherein the processor based system is configured to execute steps comprising identifying one or more groups of
15 objects in a frame of content, wherein each group of objects includes one or more shadow casters and corresponding shadow receivers, for visible shadow receivers in the identified groups of objects, establishing depth information in a small white buffer, for each of the identified groups of objects, rendering the shadow casters into a single depth map, blurring the small white buffer and upsampling the small white
20 buffer into a full sized white buffer, wherein the upsampling comprises using depth information as a metric.

BRIEF DESCRIPTION OF THE DRAWINGS

The above and other aspects, features and advantages of several embodiments
25 of the present invention will be more apparent from the following more particular description thereof, presented in conjunction with the following drawings.

FIG. 1 illustrates a simplified exemplary frame of content of a scene, according to several embodiments of the present invention.

FIG. 2 illustrates an exemplary flow diagram of a rendering flow for rendering
30 a scene according to several embodiments of the present invention.

FIG. 3 illustrates a more detailed flow diagram of the rendering flow for rendering shadows within a scene according to several embodiments of the present invention.

FIG. 4 illustrates a method for analyzing shadows accordingly to several
5 embodiments of the present invention.

FIG. 5 illustrates an exemplary 3D Grid according to several embodiments of the present invention.

FIG. 6 illustrates an exemplary 3D caster grid according to several embodiments of the present invention.

10 FIG. 7 illustrates an exemplary 3D receiver grid according to several embodiments of the present invention.

FIG. 8 illustrates an exemplary 3D intersection grid according to several embodiments of the present invention.

15 FIG. 9 illustrates a 2D screen grid being intersected with the 3D intersection grid according to several embodiments of the present invention.

FIG. 10 illustrates a flow diagram of an alternative method of rendering shadows within a scene according to several embodiments of the present invention.

20 FIGS. 11 and 12 illustrate a more detailed flow diagram of an alternative method for rendering shadows within a scene, according to several embodiments of the present invention.

FIG. 13 illustrates a small white buffer and a full sized white buffer according to several embodiments of the present invention.

FIG. 14 illustrates an apparatus or system that may be used for any such implementations.

25 Corresponding reference characters indicate corresponding components throughout the several views of the drawings. Skilled artisans will appreciate that elements in the figures are illustrated for simplicity and clarity and have not necessarily been drawn to scale. For example, the dimensions of some of the

elements in the figures may be exaggerated relative to other elements to help to improve understanding of various embodiments of the present invention. Also, common but well-understood elements that are useful or necessary in a commercially feasible embodiment are often not depicted in order to facilitate a less obstructed view of these various embodiments of the present invention.

DETAILED DESCRIPTION

The following description is not to be taken in a limiting sense, but is made merely for the purpose of describing the general principles of exemplary embodiments. The scope of the invention should be determined with reference to the claims.

Reference throughout this specification to "one embodiment," "an embodiment," or similar language means that a particular feature, structure, or characteristic described in connection with the embodiment is included in at least one embodiment of the present invention. Thus, appearances of the phrases "in one embodiment," "in an embodiment," and similar language throughout this specification may, but do not necessarily, all refer to the same embodiment.

Furthermore, the described features, structures, or characteristics of the invention may be combined in any suitable manner in one or more embodiments. In the following description, numerous specific details are provided to provide a thorough understanding of embodiments of the invention. One skilled in the relevant art will recognize, however, that the invention can be practiced without one or more of the specific details, or with other methods, components, materials, and so forth. In other instances, well-known structures, materials, or operations are not shown or described in detail to avoid obscuring aspects of the invention.

Image-based modeling and rendering methods rely on a set of two-dimensional (2D) images of a scene to generate a 3D model and then render some novel views of this scene. Such methods eliminate the need to manually create a geometric model in 3D. When an object is rendered in 3D graphics, shadows have to be rendered with respect to the placement of objects in the scene. Shadow mapping is

a process by which shadows are added to 3D computer graphics. The process of creating shadows generally involves testing whether a pixel is in the volume where any casters project over and whether the pixel is visible from the light source.

Some of the embodiments of the present invention provide one or more methods for rendering shadows in 3D computer graphics in real time that are fast and that provide sharp shadows with high resolution.

FIG. 1 illustrates a simplified exemplary frame of content of a scene. The frame 100 includes a light source 110, a shadow caster 120, and a shadow receiver 130. A shadow caster is an object that casts shadows. A shadow receiver is an object that receives shadows, meaning that shadows can be cast onto it. In one embodiment, the shadow caster 120 may be an object such as a wall or any other object, and the shadow receiver 130 may be an object such as piece of wood lying on the ground, the ground itself, or any other object. The shadow caster 120 casts a shadow onto the shadow receiver 130. The shadow caster 120 and the shadow receiver 130 comprise a group of objects that includes one or more shadow casters and corresponding shadow receivers. Stated differently, the shadow caster 120 and the shadow receiver 130 comprise a casters/receivers pair. In the following description a shadow caster is sometimes referred to as simply a “caster” and a shadow receiver is sometimes referred to as a simply a “receiver”. In some embodiments, each scene and/or portion of a scene comprises a set of casters and receivers referred to herein as a shadow set (SS).

In some embodiments, shadows are collected and computed inside the opaque phase of a rendering flow. In one embodiment, the shadows are rendered using a single shadow map. A shadow map is generated by taking a picture of a scene from the light’s point of view, also referred to as shadow’s point of view. Next inside the opaque pass each pixel is examined and transferred into shadow space and if it is determined that the pixel is further away from the light than the depth value in the shadow map, it is assumed to be in shadow.

However, having a single shadow map for a scene may raise issues with resolution according to some embodiments. That is, with certain scenes, a single map may not provide the right resolution for all pixels. For example, objects that are up

close in the scene may have too few shadow texels or shadow texture elements, and objects that are too far will have too many texels. One solution, according to some embodiments, may be to create a large map, such that enough texels can be provided for each object on the screen, i.e. maintaining at least one filtered texel per pixel, where optimal would be exactly one filtered pixel per texel. However, in certain situations, a single large map may lead to very low speeds in rendering shadows and may lead to waste of resources during the rendering process. In one embodiment, when one shadow map is used for the scene, oversampling may occur, as the shadow map contains too many samples where choosing from among the excess is wasteful. In other embodiments, undersampling may occur where the shadow map contains too few samples and extra samples must be synthesized by excess filtering and sampling. Sampling tends to be a large portion of the cost in the rendering process, and therefore, in some embodiments, it is desirable to optimize the process of shadowing to reduce oversampling and/or undersampling as much as possible. The larger the map, the more the chance is for oversampling, and the smaller the map, the more undersampling problems arise.

Accordingly, in some embodiments, the present system provides a method of rendering shadows that uses more than one shadow depth map. For example, in one embodiment, the present system employs cascade shadow maps or parallel split shadow maps. For example, in one embodiment, cascade shadow maps may be used to address undersampling and oversampling issues. Cascade maps employ the concept of splitting the viewing frustum into "cascades". According to one or more embodiments, each cascade has its own near and far plane parallel to the original near and far plane of the viewing frustum. In one embodiment, cascades are achieved by dividing the scene along the camera view direction to achieve multiple frustums each having a near and far plane.

Next, each cascade has to be rendered. That is, in one embodiment, instead of rendering a shadow map once for the entire scene, the system will have to render once for each of the cascades. To achieve this, in some embodiments, the shadow map is divided into multiple shadow maps. That is, in one embodiment, for example, the shadow map is divided into as many maps as there are cascades. For example, in one

embodiment, the scene is divided into four cascades, and according the original 1024 x 1024 shadow map is divided into four 512 x 512 maps one for each cascade. This improves the texel to pixel ratio, because the volume of the near cascades are smaller and therefore, provide more sample texels for the close up objects in the scene, and
5 provide for improved resolution.

In some embodiments, to address ALU load caused by creating cascades, the renderer will be able to specify how many cascades will be used for a particular scene. That is, in one embodiment, each scene may have a unique number of cascades depending on the resolution demands for that scene.

10 In additional or alternative embodiments, the issue of ALU load may be addressed by deferring shadows. In some embodiments, PCF sampling may be performed for shadowed geometry that is eventually occluded leading to wasted sampling in this type of opaque “overdraw”. This issue is addressed, in one embodiment, by deferring shadows. The deferred shadow is referred to herein as a
15 white buffer (WB), according to some embodiments of the present invention. In one embodiment, the white buffer is a full screen buffer, having all the results of all of the PCF sampling for a scene and/or cascade. In some embodiments, the white buffer is initially set to white, representing a maximal, or unshadowed value. In one embodiment, a full screen buffer may be used having all shadow data. In one
20 embodiment, the white buffer may be composed of a depth buffer and a colored buffer. In one embodiment, the depth buffer contents may be shared with the rest of the render pipeline, and in particular with the opaque pass. In some embodiments, the white buffer may be written to at full resolution. In one embodiment, the white buffer may be cleared and the depth may be set to its default value, and the color set to full
25 intensity, i.e. white.

In some embodiment, by using a white buffer or deferred shadow buffer, shadow transformation and PCF sampling may no longer be performed within the opaque pass. In such embodiments, instead, shadowing in the opaque pass may comprise point-sampling a screen sized texture, i.e., the white buffer. In such
30 embodiment, the process of determining per pixel shadows is done in a single step, instead of having to potentially sample multiple times. In one embodiment, the white

buffer is both overdraw independent and shadow sampler independent. That is, in some embodiments, the opaque shader samples a single map. Further, in some embodiments, there are four channels available, therefore, in such embodiments, up to four light source's shadow data may be put into a single deferred map.

5 In some embodiments, deferring shadows further allows for having freedom with respect to the size of maps for each cascade of a scene. That is, in some embodiments, any map size up to the original map size may be chosen for each of the cascades in the scene. In one embodiment, this is possible because more than one shadow map may be successively put onto the actual map before being accumulated
10 into the WB.

 In some embodiments, using deferred maps, the process of shadow rendering may be further optimized by using depth values obtained from the z-pass, e.g. a ZCull unit. In one embodiment, a ZCull unit provides both a Zfar and Znear. In such
15 embodiments, the system can obtain data to exclude pixels having a depth that is farther away than what is already in the depth map. Further, having a Znear, in some embodiments, allows ignoring pixels independent of their incoming interpolated depth. That is, in one or more embodiments, the ZCull unit may enable the system to wholesale reject pixels if the depth underneath the pixels is outside of a selected
range.

20 In some embodiments, for each cascade a full screen pass is performed. However, in some embodiments, using Zcull depth bounds, only appropriate pixels are drawn during the pass, such that only pixels in the specific cascade are drawn for the specific white buffer.

 In some embodiments, while cascades and deferred shadow maps provide
25 larger map sizes, the map size may still be limiting for certain scenes. That is, in some embodiments, a ratio of one texel per screen pixel, or even better ratio may be desired. Having an upper bound of the shadow map size, limited by the original map size, leads to not being able to achieve such ratios for certain scenes. In one embodiment, this is addressed with the concept of tiling the shadow map.

30 In one embodiment, the depth map may be tiled, in 2D. That is, in some

embodiments, after the viewing frustum is tiled by cascades along the volume, the shadow map is further tiled in 2D. In one embodiment, the shadow maps may be tiled as many times as desired and accumulated one at a time into the white buffer. In some embodiments, MIN blending may be used to achieve this.

5 In one embodiment, tiling may be performed on all cascades of the scene. In another embodiment, tiling may be performed on less than all cascades. For example, in one embodiment, only the first few cascades may be tiled and the other cascades may be rendered without tiling. In some embodiments, additional cascades may be added. In one embodiment, scenes may be rendered with better quality by
10 adding additional cascades.

 In one embodiment, the setting for each cascade may be user, e.g. game maker, etc., controllable. In one embodiment, the user's view of the shadows may involve multiple explicit cascades, and settings for each cascade, such as resolution, i.e., texels per world unit, may be set for each cascade by the user. In some
15 embodiments, further, the user may be provided with a "quality" knob to control the amount of PCF sampling performed per cascade.

 The process of optimizing the white buffer and ultimately shadow rendering, according to several embodiments of the present invention is based on the concept of shadows and how objects receive shadows. Namely, one or more embodiments of the
20 present invention take advantage of determining when pixels should or should not receive shadows. According to several embodiments, certain pixels do not receive shadow and therefore may be occluded to optimize the shadow rendering process. Accordingly, in several embodiments, pixels are tested to determine whether they are within a volume of space that one or more casters project over. In further
25 embodiments, additionally or alternatively, for pixels within a volume where casters project over, it is determined whether the caster is supposed to affect the pixel.

 FIG. 2 illustrates an exemplary flow diagram of a rendering flow 200 for rendering a scene according to several embodiments.

 In step 210 a Zprepass is performed on all rendering geometry in the scene.
30 Next, in step 220, the scene is divided into one or more cascades. Next, in step 230,

the first cascade is analyzed for shadow information. Next, in step 240, a white buffer is generated for the cascade. In step 250 it is determined whether a next cascade is present, and when it is determined that a next cascade is present the above step 230 and 240 are repeated for the next cascade. If on the other hand, if in step 250 it is
5 determined that all cascades of step 220 have been analyzed and a white buffer has been generated for all cascades, then the process continues to the opaque stage of the rendering flow.

FIG. 3 illustrates a more detailed flow diagram of the rendering flow for rendering shadows within a scene according to several embodiments.

10 As illustrated, in step 310, the system performs a Zprepass for all rendering geometry in a scene. In one embodiment, the depth values may be later used to reconstruct receivers. In some embodiments, during this step a “type” is determined for each receiver. In one embodiment, a value indicating the determined “type” for the receivers is written in the stencil. For example, in one embodiment the “type”
15 value may represent one of: “none”, i.e., no type available for the receiver, “background”, “minor”, or any other identifier. In one embodiment, the remaining stencil bit may be used to write a more geometrically unique identifier for each of the receivers.

Next, in step 320, the white buffer is cleared. As discussed above, white
20 buffer (WB), may be a full screen buffer that has all shadow data. In one embodiment, the white buffer is a full screen buffer identical in size to the back buffer. In one embodiment, the white buffer may be a 1280x720 buffer. However, it should be obvious to one skilled in the art that different size buffers may be used for different purposes. In one embodiment, the white buffer may be anti-aliased. In some
25 embodiments, the white buffer may be composed of a depth buffer and a color buffer. In some embodiments, the white buffer is written to at full resolution. In some embodiments, the white buffer is initially cleared and the depth is set to its default value, and the color is initially set to full intensity (white).

Next, in step 330, the scene is divided into a number of cascades. In different
30 embodiments different number of cascades may be used according to several factors.

For example, in one embodiment, the scene may be divided into three cascades. In some embodiments, for each scene, the number of cascades may be varied based on several factors including for example the requirements of a scene and/or the preferences of the user, e.g. game maker.

5 In step 340, casters are divided into categories. For example, in one embodiment, the categories may include: (a) casters that cast on everything, (b) casters that cast on everything but don't otherwise render, (c) casters that cast only on background, and not on themselves, (d) casters that cast only on not background but don't otherwise cast. Other categories may be used and the above mentioned
10 categories only provide examples of possible categories for casters.

 Next, in step 350, the process determines the light sources in the scene. In some embodiments, each light source within the scene may cause the casters to set a unique set of shadows on the receivers, and therefore, the system may analyze the light sources separately. In one embodiment, the shadow information for each light
15 source may be written into a different channel within the white buffer. In another embodiment, all shadow information may be written into the same channel in the while buffer.

 In step 360, for the first cascade, the process analyzes the shadows with respect to each category of caster, for each light source. In one embodiment, this step
20 comprises constructing a 3D caster grid, constructing a 3D receiver grid, intersecting the grids to form a new grid, identifying casters that intersect this new intersection grid and using the new grid to exclude portions of the screen from receiving shadows. The process for analyzing shadow according to some embodiments is described in further detail below in FIG. 4.

25 In step 370, the process renders the shadow map for all casters found to cast a shadow as determined during the analyzing step 360. In one embodiment, these casters refer to those that intersect with the receivers in the caster receiver grid of step 460 of FIG. 4. In one embodiment, the casters for each light source, category, cascade and tile are rendered into the shadow map. Further, in some embodiments,
30 when rendering 2D cells for drawing the receiver pixels, the receiver sections

determined in step 440 of FIG. 4 are drawn. In a further embodiment, during the rendering of the shadow map, using the depth information available, e.g. as determined in step 410, the pixels whose z buffers fall outside of the depth bounds may be ignored and not drawn. In such embodiments, the rendering of the shadow map is optimized as many pixels can be thrown away.

In one embodiment, before rendering the shadow map, the shadow map is tiled. In some embodiments, the tiling of the shadow map may increase the rendered resolution, which relates to the shadow texel to pixel ratio. In one embodiment a setting relating to a desired shadow texel to meter ratio is provided. In some embodiments, to keep a constant number of shadow texels per world unit, it may be necessary to make some number of sub-cascades, which uses all the information gathered above. In one embodiment, the process may breakup the number of shadow maps assigned to a cascade, i.e. tile the shadow map. In one embodiment, any number of tiles per cascade may be created to satisfy the required shadow texel to world unit requirement.

In one embodiment, at runtime a maximum and/or minimum number may be set for the number of tiles. In one embodiment, for example the minimum number of tiles may be set equal to 1. In some embodiments, the maximum number of tiles may be set to any number deemed appropriate for the specific situation and what is desired. Further, in some embodiments, to further optimize the process, tiles with very small dimensions on a side may be omitted by slightly adjusting the shadow texel to world unit ratio as needed. In one embodiment, tiles may be trimmed when they occupy a very small number of texels and/or the tile is a direct result of overflowing of a previous tile's texels.

Next, the process moves to step 380 and constructs a white buffer based on the shadow map of step 370. In some embodiments, a shadow map is rendered for each unique cascade, light source and category combination. In such embodiments, the system will repeat the steps 360 and 370 for each category and for each light source within the cascade. For example in one embodiment, the optional steps 372 and 374 may be performed to ensure that a shadow map is rendered for each unique combination. Next, if it is determined that a shadow map is rendered for each

combination of light source and category for the cascade, the process continues to step 380 where the white buffer is generated for the cascade.

In one embodiment, for each light source, category, cascade and tile combination, the depth buffer is sampled during this step. In some embodiments, additionally or alternatively, a camera-space or world space position is determined and reconstructed. The camera space or world space position is then transformed into texture coordinates according to one embodiment. In a further embodiment, the shadow map is then sampled. In one embodiment, as a quality improvement measure multiple PCF samples may be taken. According to one embodiment, the process of creating a shadow map is repeated for each light source and category combination for each cascade.

In one embodiment, for each 2D grid cell, a stencil value of from the Zprepass of step 310 may be used to discriminate which pixels receive shadow. For example, in some embodiments, if the stencil indicates that a pixel is a background pixel, but the category is “cast only on not background, but otherwise don’t render,” from step 320, then that pixel will be ignored by the stencil hardware. In some embodiments, 3D extent information gathered with respect to each 2D grid is used to perform depth bounds testing to eliminate un-shadowed pixels.

In several embodiments, the ZCull unit includes one or more features that are effective in improving the performance of the shadowing process. In some embodiments, in addition the Znear and Zfar, the system further includes a stencil reject buffer. In one embodiment, geometry culling is performed per draw call bounds. In some embodiments, accordingly, the system may be capable of rejecting on the basis of depth and/or stencil. In some embodiments, the rejection can be done on the basis of both depth and stencil at the same time.

In one embodiment, rejection may be very fast. For example, in one embodiment, rejection may be approximately eight times faster than the nominal accept rate. In one embodiment, the nominal accept rate corresponds to running a one cycle pixel shader. In some embodiments, the shadow shader may comprise 5 cycles. In such embodiments, this will lead to the process being approximately 40 times faster. In one embodiment, using the stencil minor characters are drawn with better

precision and faster speed.

In one embodiment, a resultant shadow value is accumulated into the WB, using MIN blending. In one embodiment all shadow values may be accumulated into the same color component. In an alternative embodiment, the shadow values are
5 segregated by light source, so that each lights gets its own set of shadow samples, or shadow set. In one embodiment, using MIN blending the order of operations for laying down shadow samples may be ignored.

In one embodiment, at the end of step 385, the WB includes all the shadow information. In some embodiments, where Multi-Sample Anti-Aliasing (MSAA) is
10 employed, the WB will be at MSAA resolution in number of pixels, but in terms of data it will only be at non-MSAA effective resolution. In some embodiments, this can cause artifacts when the lighting pass runs, since the lighting pass samples will not always be in agreement with the shadow samples that exist in the WB. In some
15 embodiments, as described above, stencil values written in step 310 include a semi-unique group ID for each group of geometry. In some embodiments, the stencil values are used to dilate the WB. In some embodiments, for example, a “plus” pattern of MIN of four neighbors of every pixel MINed with each WB pixel may be used in dilating the WB. This may be useful in hiding nearly all depth discontinuity artifacts. In some embodiments, the dilation may involve examining the anti-aliasing pixel
20 transition between light and dark and growing the dark based on the assumption that dark provides a better picture.

Next, in step 385, it is determined if a WB has been created for all of the cascades. If it is determined that a WB has not been constructed for one or more of the cascades, the process returns to step 340 and repeats steps 340-380 for each
25 cascade generated in step 330.

On the other hand, if in step 385 it is determined that a WB has been constructed for each of the cascades of step 330, the process continues to step 390 and renders the receiver geometry. In some embodiment, when rendering the
opaque/lighting pass, the results of the analyzing of the shadow receivers may be used
30 to determine the type of pixel shader in the opaque pass. For example, in some embodiments, marked geometry, i.e., receiver geometry that intersects one or more

casters, chooses a pixel shader that samples a WB. For other potentially receiving geometry, e.g. geometry not marked during the shadow analyzing process 400, the geometry is not shadowed.

5 In one or more embodiments, baked shadow integration may be performed according to several methods. In one embodiment, baked shadow may take the form of HDR light maps, HDR vertex based light, etc. These forms of the baked shadow may contribute to the “direct” portion of the lighting equation, to contrast with indirect, bounced or ambient light. In some embodiments, an artist may choose which lights should be shadowable. In some embodiments, indirect light is not shadowed.

10 In one embodiment, certain receivers within the scene may be shadowed by hidden casters. According to some embodiments, hidden casters are rendered into the shadow map, but do not exist as visual geometry. In one embodiment, some hidden casters may only render on pixels that do not receive baked shadows. In some embodiments, such casters are not included in the same shadow map as those casters that cast on everything. In one embodiment, the hidden casters have their own shadow map and therefore go through a separate set of full screen passes, independent of objects that cast on everything.

FIG. 4 illustrates a method 400 for analyzing shadows accordingly to several embodiments of the present invention. In step 410, a 3D caster grid is constructed. In one embodiment, the caster grid constructed in step 410 comprises a grid from the particular light source point of view that fully encloses the current cascade frustum. In some embodiments, a fixed number of cells may be used in the grid. For example, in one embodiment, an 8x8 grid may be used. In other embodiments, other grid sizes may be used.

25 In one embodiment, for each light source, for each category, for each cascade a 3D caster grid is constructed. In one embodiment, before beginning the process in step 410, for each category and caster for each light source, all possible casters and receivers are collected. In one embodiment, the set of all casters and receivers may be referred to as a shadowing set (SS). In a further embodiment, before beginning the process 400, after collection the shadowing set, for each of the casters and/or receivers within the shadowing set, the per draw call bounding information for the caster and

receivers is collected. In some embodiments, the call bounding information comprises spheres. In alternative embodiments, Axis Aligned Bounding Boxes (AABBs), Oriented Bounding Boxes (OBBs) or other types of bounding information may be used.

5 Next, in step 415, casters in the shadow set are projected onto the 3D grid using the bounding information collected, and the 3D grid is filled in. In some embodiments, each grid cell may be viewed as an AABB in light space. In one embodiment, where spheres are used as the bounding information, projecting casters will form capsule like bodies.

10 After the projection, in step 420, each intersected cell is grown to fully encompass the caster. In one embodiment the growth is limited to the actual cell bounds within grid. In one embodiment, during this step, a union of the cell's current bounds and the projected caster's bounds is performed. In such embodiment, any part of the union which is outside the cell's maximum boundary within the 3D grid is
15 clipped so that the extents of the cell do not exceed the maximum boundaries. In one embodiment, "union" is similar to an OR function in that it refers to the volume being occupied by only one.

 In one embodiment, a receiver only receives a shadow if a projected caster intersects that receiver. In some embodiments, the 3D caster grid is used to avoid
20 testing every caster individually against each receiver to determine if the caster hits the receiver. Instead, in some embodiments, the 3D caster grid may be used to optimize the test for speed.

 Next, in step 425, a second 3D grid, the receiver grid, is constructed. In one embodiment, the 3D receiver grid covers the same volume as the caster grid. In one
25 or more embodiments, the bounding information of the receivers within the shadow set is used to construct the 3D grid. In some embodiments, the receiver, unlike the caster, does not project. Therefore, non-projected bounding volumes are used for the receivers within the receiver grid. In one embodiment, to optimize the construction of the caster grid, it is not necessary to fill in cells that are not touched by any casters in
30 the caster grid.

In step 430, cell by cell, the receiver grid is intersected with the caster grid to determine if any of the receivers intersect the caster grid. If it is determined that the any receivers intersect the caster grid, in step 435, the receiver is marked for later rendering, and the process moves to step 440. In one embodiment, intersecting the
5 cells of the receiver grid with the caster grid is similar to an AND function in that it determines that the volume is shared by both. As such, in some embodiments, if a specific receiver touches only empty caster grid cells, the there is no intersection, and the specific receiver is thrown away. Accordingly, in such embodiments, receivers that do not receive shadows for the specific light source, cascade and category are
10 thrown out. Alternatively, if a specific receiver has an intersection with a caster grid cell, then that receiver is marked for rendering. In one embodiment, the information is used to develop a list of receivers for later rendering into a shadow map. In one embodiment, for example, receivers are marked for to use the shadowing pixel shader in the opaque pass if it is determined that the receiver intersects with the caster grid
15 cell for any of the cascades, even if only one cascade. In some embodiments, those receivers that are not marked during this step may be assumed to not be shadowed by the casters and therefore may choose a pixel shader that does not sample the WB.

In step 440 the cells covered by each receiver that intersects the caster grid as determined in step 430, are grown. After completing step 440, i.e., after all of the
20 receivers within the receiver grid have been intersected with the caster grid and the intersecting receivers have been marked for rendering and the respective cells covering the receivers have been grown, the process then continues onto step 450.

In step 450, according to some embodiments, the caster grid and the receiver grid are intersected cell by cell to construct a new intersection grid, which is the
25 intersection of the caster and receiver grids. In one embodiment, the 3D intersection grid covers the same volume as the caster grid and receiver grid. In one embodiment, in step 455 it is determined if any casters intersect this new final intersection grid.

If in step 455, if it is determined that a caster intersects the new final intersection grid, the caster is marked for rendering in step 460. In one embodiment, a
30 caster may be marked for later rendering if there is any intersection at all. In some embodiments, if a particular caster does not intersect the final intersection grid, it may

be concluded that the caster does not touch any receivers and therefore the caster may be thrown out. Accordingly, in some embodiments, in step 460, a list of casters for later rendering into shadow map is developed.

Next, in step 465, the screen is apportioned into a number of 2D grids. In some embodiments, an 8x8 grid may be used, however it should be understood that any size grid may be used, for example, based on system or project requirement. The 2D screen space is then projected in 3D, which represents the volume of space that the 2D grid cells project over based on the current view projection being employed. In one embodiment, the 2D grid is projected such that each cell of the 2D grid forms a unique 3D frustum.

Next, in step 470, according to several embodiments, for each screen grid cell projected into a unique frustum, it is determined if there are any 3D cascade cells within that frustum. That is, it is determined if there are any intersections. In one embodiment, if it is determined that there are no 3D cascade cells within the frustum, i.e., screen grid, then that 2D cell is determined as not receiving any shadows for the current light source, category and cascade combination. In one embodiment, upon determining that there is no intersection, the process moves to step 475 and tests the next cell within the grid, such that the process is repeated for all cells of the 2D screen grid.

If it is determined that there are intersections, i.e. there are 3D cascade cells within the frustum, then in step 480, a 3D AABB intersection volume is determined. In some embodiments, depth information may be determined from the 3D grid. This information may be used, according to some embodiments, by the hardware (HW) depth bounds functionality to exclude pixels based on their z-buffer depths. Accordingly, the 3D grid is used to both narrow the 2D area that receives shadows as well as narrow the depth of the corresponding 3D volume. In several embodiments, accordingly, only pixels that are in the identified 2D areas and within the depth range are drawn. In one embodiment, this leads to many pixels being thrown away.

FIG. 5 illustrates an exemplary 3D grid or shadow volume 500, enclosing the cascade frustum 520. A cascade near plane 522 and far plane 524 are illustrated to

define the boundaries of the cascade frustum 520. In this exemplary embodiment, the direction of the light source is illustrated by an arrow 530. As seen, in this exemplary embodiment, for demonstration purposes, the light source direction is straight down. The shadow volume is partitioned into a grid. In this exemplary embodiment, the shadow volume 500 is partitioned and the shadow's X and Y coordinates are in the same plane as the near plane of the 3D shadow volume, and Z is moving away from the shadow volume near plane, 3D box illustrated in FIG. 5. As the shadow volume covers at least the portion of the scene encompassing a cascade, for illustrative purposes, it is assumed that the scene comprises both casters and receivers, such that there are both casters and receivers inside this volume. The 3D volume is analyzed by the process 400 to determine which casters cast shadows on which receivers from the scene. In one embodiment, after this information is gathered, the process 400 then moves to a 2D screen grid to determine how everything should appear in the viewer's perspective.

FIG. 6 illustrates an exemplary caster grid as seen from the side. For illustration purposes the frustum is not shown, however, the frustum if shown would be on the right directed towards the left. In one embodiment the caster grid covers the same volume as the 3D grid 500 of FIG. 5. As seen in FIG. 6 the particular grid only comprises one caster 610, according to the shadow set information. The caster is illustrated as having been placed in the grid according to the bounding information and projected. The intersecting cells of the grid have been grown. In one embodiment, the caster grid of FIG. 6 is a result of steps 410-420 of process 400.

FIG. 7 illustrates an exemplary receiver grid according to one or more embodiments of the present invention from the side. The 3D grid of FIG. 7 covers the same 3D volume as the 3D caster grid 500. In one embodiment, the receiver grid covers the shadow volume 500 illustrated in FIG. 5. FIG. 7 illustrates 3 receivers 710, 720 and 730 within the 3D grid, according to information from the shadowing set gathered. In some embodiments, at render time the receivers only exist as pixels in the zbuffer, however, according to one embodiment, the system will have bounding information available for the receivers. In one embodiment, the receivers do not project.

FIG. 8 illustrates the intersecting the receivers with the projected casters. As is seen, receivers 710 and 720 intersect with the caster 610, but receiver 730 does not intersect with any casters. In some embodiments, the system generates an AABB of the intersection of what the caster covers and the receivers, i.e. boxes 810 and 820.

5 FIG. 9 illustrates the 2D grid being intersected with the 3D intersection grid. The lower lines 910 and 920 of FIG. 9 represent the projection of a 2D screen space cell in 3D. In one embodiment, the space between lines 910 and 920 represents the volume of space that the 2D cell can see within the current cascade. According to some embodiments, each cell of the 3D intersection grid is tested against this space.

10 In some embodiments, this comprises intersecting the 3D grid with the 2D projected frustum. Further, based on the intersection, the depth bounds for the particular 2D cell may be further determined. For example, as seen in FIG. 9, the number of pixels that have to be drawn is reduced by more than half at least along the Y direction as can be seen through the side view of the 3D grid. That is, the process may determine

15 where the 3D cell AABBs intersect the 2D projected frustum. In this exemplary embodiment, the 2D cell area to be rendered is reduced from the area covered by lines 910 to 920 to the area between lines 915 and 920. Further, by adjusting the depth bounds to the bounds of the volume at which the 3D grid and the 2D projection intersect, the depth bounds are reduced from the entire length of the frustum, i.e. 940

20 and 950, to the new bounds 960 and 970.

The above methods and processes provides an optimized rendering flow with a shadow rendering process where the number of passes and sampling is reduced, and therefore shadows are provided at an optimal resolution while maintaining a desirable rendering speed.

25 FIGS. 10-13 describe an alternative process for rendering shadows. Specifically, some of the embodiments of the present invention provide one or more methods for rendering shadows in 3D computer graphics in a manner that reduces aliasing artifacts and that lends a soft appearance to the shadows. For example, in some embodiments the present invention provides a series of methods for rendering

30 image based projected shadows with multiple depth aware blurs, a purpose of which includes reducing aliasing artifacts and lending a soft appearance to shadows rendered

into a multisampled anti-aliased buffer.

Referring to FIG. 10, there is illustrated an example of a method 1000 in accordance with an embodiment of the present invention. The method 1000 may be used in 3D computer graphics for rendering image based projected shadows with multiple depth aware blurs. In some embodiments, the method 1000 may be used for rendering such shadows in real time. The method 1000 may be applied to any type of content, such as for example video games, both online and offline game types, simulations, other types of videos, etc.

The method 1000 begins in step 1010 where a frame of the content is analyzed. Specifically, one or more groups of objects in the frame are identified such that each group of objects includes one or more shadow casters and corresponding shadow receivers. In one embodiment for example, the frame of content is similar to frame 100 illustrated in FIG. 1.

In step 1020, depth information for visible shadow receivers in the identified groups of objects is established in a small white buffer. In some embodiments, this buffer is actually composed of a depth buffer and a color buffer. The small white buffer may be referred to as a deferred buffer. In some embodiments, this step may be performed by rendering a depth prepass to the small white buffer. In some embodiments, however, it is not necessary to render receivers into the white buffer if a depth prepass has already been performed at the full resolution. In this case, the small white buffer's depth may be constructed by downsampling the larger depth buffer. Point sampling the large buffer is sufficient.

In step 1030, for each of the identified groups of casters/receivers pairs, the shadow casters are rendered into a single depth map, and if necessary, the shadow receivers are rendered to the small white buffer.

In step 1040, the small white buffer is blurred. In some embodiments a low resolution blur may be used.

In step 1050, the small white buffer is upsampled into a full sized white buffer. That is, the small white buffer is made bigger. In some embodiments, the upsampling technique may use depth values or information as a metric. That is, in some embodiments the upsampling comprises using depth information as a metric. A technique for performing the upsampling that may be used in some embodiments will

be described below. The upsampling technique described below upsamples the deferred buffer in a manner that is can preserve more detail than simply using a box filter. In some embodiments, the upsampling technique described below may comprise features of bilateral upsampling techniques.

5 In step 1060, the shadow receivers in the identified groups of objects are rendered by sampling the upsampled full sized white buffer.

 An embodiment will now be described in greater detail with reference to FIGS. 11 and 12. Specifically, there is a method 1100 in accordance with an embodiment of the present invention. The method 1100 may be used in 3D computer
10 graphics for rendering image based projected shadows with multiple depth aware blurs, which helps to reduce aliasing artifacts and lend a soft appearance to shadows rendered into a multisampled anti-aliased buffer. In some embodiments, the method 1100 may be used for rendering such shadows in real time. The method 1100 may be applied to any type of content, such as for example video games, both online and
15 offline game types, simulations, other types of videos, etc.

 In step 1110, depth is cleared and a depth prepass is performed with all opaque and cutout geometry in the content.

 In step 1120, for every frame of the content, all objects or subsets of objects are gathered into collections of shadow casters/shadow receivers pairs for those active
20 shadowing lights affecting visible receiver geometry. Thus, the shadow casters/receivers pairs of objects in the frame include one or more shadow casters and corresponding shadow receivers. In the discussion herein, the term “shadow casters” may be simplified to “casters”, and the term “shadow receivers” may be simplified to “receivers”.

25 In step 1130, a “white buffer” is cleared and the depth is set to its default value, and the color is set to full intensity (white). This buffer is referred to herein as the full sized white buffer.

 In step 1140, only for all visible receivers selected previously, another depth prepass is rendered, but to a much smaller “white” buffer. This buffer is referred to
30 herein as the small white buffer. Alternatively, it is also sufficient to downsample the depths from the full sized white buffer.

 Next, for each casters/receivers pair, the casters for the pair are rendered into a

single depth map. Then, each set of receivers in the pair are rendered to the small white buffer. Alternatively, downsampled depth values from the small buffer may be used. These steps are illustrated by steps 1150, 1160 and 1170. Namely, in step 1150 the casters in a first casters/receivers pair are rendered into a single depth map. Then
5 in step 1160, each set of receivers in the pair are rendered to the small white buffer.

In some embodiments, step 1160 of rendering each set of receivers in the pair to the small white buffer may be performed using an appropriate amount of filtering samples. By way of example, in some embodiments Percentage Closer Filtering (PCF) may be used. The sample width may be adjusted based on the derivative of
10 the depth as well as the derivatives of the shadow maps sampling texture coordinates. Optionally, in some embodiments, a Poisson sampling distribution may be used to mask high frequency “steps” resulting from undersampling the caster depth map. In some embodiments, the shadowed sample may be written into the small white buffer using a MIN blend mode.

15 In step 1170, the system determines whether or not there are any more casters/receivers pairs. If so, steps 1150 and 1160 are repeated. Steps 1150 and 1150 are repeated for each casters/receivers pairs. Once step 1170 determines that there are no more casters/receivers pairs, then the process flow proceeds through connector 1210 and on to step 1220 (FIG. 12).

20 In step 1220, the small white buffer is blurred. In some embodiments, the small white buffer may be blurred with a depth aware low pass filter. In some embodiments, for those samples that are too far in terms of depth distance from the sample point, their influence is not included in the blur. That is, in some embodiments, the blurring comprises not including an influence of samples having a
25 depth that is beyond a certain distance from a sample point. In some embodiments, for those samples within an acceptable depth distance, they are weighted based on distance and include in the filter. That is, in some embodiments, the blurring comprises weighting based on distance each sample having a depth that is within a certain distance from a sample point. Again, the word “far” refers to depth distance.

30 Blurring the small white buffer is effectively performing a low resolution blur. The small white buffer is blurred with itself, and in some embodiments the blur is based on the relative depth values of the pixels. That is, the Z values (or depth values)

are taken into account in performing the blur. In some embodiments, an in-place blur is performed by deciding which pixels to blur based on their relative depth values. Thus, depth information or values are used as a metric. In some embodiments, the blur performed on the small white buffer may be a bilateral blur.

5 Taking the Z value differences between the pixels in the small white buffer into account helps to preserve detail. This is because the Z value differences are used to determine how much influence a neighboring pixel will provide in the blur, that is how much of a neighboring pixel to add in. In some embodiments, if the Z values of neighboring pixels are similar (i.e. the Z value difference is small), then the
10 neighboring pixels will be used to provide an influence in the blur. But if the Z values of neighboring pixels are very different (i.e. the Z value difference is large), then the influence of the neighboring pixel will not be used in the blur. The Z value difference may be large if, for example, a shadow boundary lies between the two pixels. In this case, not using the influence of the pixel in the blur helps to preserve detail because
15 the shadow boundary will remain sharp and not be smudged by the blur. Performing an in-place blur on the small white buffer helps to get smooth changes in shadows. Thus, using depth information or values as a metric helps to determine how much influence neighboring pixels will have in the blur. In some embodiment, pixels are effectively thrown out that are too different in depth. This way, pixels are only
20 smudged together when it makes sense to smudge them together. In some embodiments, the blur may be a bilateral blur.

 In step 1230, the small white buffer is upsampled into the full sized white buffer. In some embodiments, the small white buffer may be upsampled into the full sized white buffer using the high resolution depth information stored in the first depth
25 prepass (e.g. step 1110) to weight each of a number of “bilinear” low resolution samples. In some embodiments, the small white buffer may be upsampled into the full sized white buffer using the high resolution depth information stored in the first depth prepass (e.g. step 1110) to weight each of, for example, four “bilinear” low resolution samples.

30 Similar to the previous blur in step 1220, in some embodiments of step 1230, the weights are set to zero for each sample whose depth is too far from the high resolution sample point’s depth. That is, in some embodiments, the upsampling

comprises setting weights to zero for each sample having a depth that is beyond a certain distance from a high resolution sample point's depth. In some embodiments, for those samples within an acceptable distance, they are weighted based on distance and included in the filter. That is, in some embodiments, the upsampling comprises
5 weighting based on distance each sample having a depth that is within a certain distance from a high resolution sample point's depth. In some embodiments, if no samples are close enough, then only the bilinear weight is used. In some embodiments, this kind of filter may include techniques used in a bilateral filter. In some embodiments, the upsampling may comprise bilateral upsampling.

10 An example embodiment for step 1230 of performing the upsampling will be further described below.

In step 1240, the receiver geometry is rendered for lighting purposes. In some embodiments, the receiver geometry is rendered for lighting purposes by sampling the now upsampled high resolution white buffer. That is, in some embodiments, the
15 shadow receivers are rendered by sampling the upsampled full sized white buffer.

In some embodiments, step 1240 may be performed by using multisample anti-aliasing (MSAA). If using MSAA, fewer samples will run than exist in the high resolution white buffer. In this case, in some embodiments, the instant depth may be taken and compared to the depths contained in the high resolution white buffer. If all
20 samples are roughly the same distance from the instant depth, then all high resolution white buffer samples may be evenly weighted and the aggregate shadow value used to attenuate (shadow) the contribution of any dynamic lights that are marked to gather shadows from the white buffer. In some embodiments, any samples whose depths are too far from the closest white buffer depth are discarded, and the remaining samples
25 are used to form an aggregate shadow value for attenuation purposes.

In step 1250, the rendered shadow receivers may be displayed on a display.

An example embodiment of step 1230 in which the small white buffer is upsampled into the full sized white buffer will now be described. FIG. 13 illustrates a small white buffer 1310 that is to be upsampled into a full sized white buffer 1320.
30 As mentioned above, the small white buffer 1310 may be referred to as a deferred buffer. The process of upsampling is sometimes referred to as "zooming up" the small white buffer 1310 to the full sized white buffer 1320, or making the small buffer

into a big buffer. Portion of pixels of small buffer map to pixels in big buffer. In this example, the upsampling is going from 2 x 2 (i.e. 4 pixels) to 4 x 4 (i.e. 16 pixels). It should be understood, however, that the methods and techniques described herein may be applied to upsample any size small buffer to any size full size buffer.

5 One way to perform upsampling is to use a box filter to zoom up the small white buffer 1310. With this technique, for each pixel in the full sized white buffer 1320, weights are assigned to the pixels in the small white buffer 1310. The weights assigned to the pixels in the small white buffer 1310 determine the amount of influence that the those pixels will have on the particular pixel in the full sized white
10 buffer 1320. The weights assigned to the pixels in the small white buffer 1310 will typically sum up to one.

 One problem with the box filter technique is that it deals only with the color of the pixels and does not take into account the depth values of the pixels. This can lead to the undesirable result of a pixel in the small white buffer 1310 providing color
15 influence when it really should not because it is at a much different depth than the pixel in the full sized white buffer 404 that it is influencing.

 For example, the cross-hatched pixels 1360, 1362, 1364, 1366, 1368, 1370, 1372, 1374, 1376, and 1378 in the full sized white buffer 1320 may represent a person's arm. Similarly, the cross-hatched pixel 1338 in the small white buffer 402
20 represents the person's arm at a lower resolution. The person's arm may be close to the camera, and so the Z value (i.e. the depth) of those pixels may be approximately $Z = 0.1$ (where the camera is $Z = 0.0$). The remaining pixels 1342, 1344, 1346, 1348, 1350, and 1352 in the full sized white buffer 1320, and the remaining pixels 1332, 1334 and 1336 in the small white buffer 1310, may represent the floor. The floor is
25 farther away (more distant) from the camera, and so the Z value (i.e. the depth) of those pixels may be approximately $Z = 0.4$.

 The box filter technique would cause pixel 1334 in the small white buffer 1310 to influence the color of pixel 1370 in the full sized white buffer 1320. However, this is undesirable because pixel 1334 is clearly the floor, and pixel 1370 is
30 clearly the person's arm. Similarly, the box filter technique would cause pixel 1338 in the small white buffer 1310 to influence the color of pixel 1348 in the full sized white buffer 1320. However, this is undesirable because pixel 1338 is clearly the

person's arm, and pixel 1348 is clearly the floor.

In order to avoid these disadvantages, embodiments of the present invention use depth (i.e. Z values) as a metric. That is, in some embodiments, the metric that is used is depth. For example, in some embodiments, pixels in the small white buffer
5 1310 are not used to influence pixels in the full sized white buffer 1320 that have a much different Z value. Namely, depth information is used to throw out certain pixel weights because they have bad information. The pixel weights to throw out are identified based on depth information (i.e. Z values). In some embodiments, the depth values of the pixels in the full sized white buffer 1320 may be determined by
10 taking the proportionate depth amount of the pixels in the small white buffer 1310. That is, for each of the sixteen pixels in the full sized white buffer 1320, take the proportionate depth amount of the four pixels in the small white buffer 1310.

In some embodiments of the present invention, the step of upsampling generally involves looking at the depth value, constructing what the weights would
15 be, and then throwing out values or re-weighting values based on differences in depth. Thus, this is a conditional process because it throws out some values. A determination is made for each pixel as to whether the high resolution depth matches the low resolution depth. If not, the influence of the pixel is thrown out or reweighted.

For example, in some embodiments, the step of upsampling may involve
20 constructing what the pixel weights for the pixels in the small white buffer 1310 would be for a particular pixel in the full sized white buffer 1320. Then, the depth values (i.e. Z values) for the pixels in the small white buffer 1310 are compared to the depth value of the particular pixel in the full sized white buffer 404. This essentially
25 determines whether the high resolution depth matches the low resolution depth. For pixels in the small white buffer 1310 where the difference in the depth value is large, the influence of those pixels is either thrown out or re-weighted. This is because those pixels would contribute bad shadow information. Not using the influence of those pixels helps keep the shadow edge straight. Namely, this preserves sharp edges
30 by systematically excluding pixels across the discontinuity from consideration.

Thus, in some embodiments, the pixel weight information is used in the upsampling, unless it is found that the depths vary by more than a certain amount.

For pixels where the depths vary by more than a certain amount, the weight information for those pixels is either thrown out, or they are reweighted. Thus, in some embodiments, pixels are thrown out having depths that do not match. That is, they are thrown out based on their Z value. Not using the influence of such pixels, or
5 throwing them out, helps to preserve detail. Detail is preserved because the color of a pixel in the small white buffer 1310 is not used to influence the color of a pixel in the full sized white buffer 1320 having a much different depth. In this way, depth information or values (i.e. Z values) are used as a metric. It is noted that the blurring step described above is similar in that depth is also used as a metric.

10 In some embodiments, the weight of the low resolution samples may be determined as follows. Specifically, a range of Z values may be chosen. That is, a delta Z between the small and large buffer may be chosen to define a tolerance. Pixels that are within the tolerance of Z values may be treated as the same. That is, even if they do not match, they are treated as the same if they are close. A small
15 tolerance may give more detail. A small tolerance may avoid it leaking to something else. In some embodiments, the artist may have control over the tolerance. That is, the artist can change/modify the tolerance. Thus, the tolerance may be used to identify sameness. For pixels in the tolerance, they may be treated the same. In some embodiments, edge pixels may be faded off slowly.

20 Thus, the upsampling process involves zooming up a small buffer to a big buffer. In some embodiments, the metric of correction is depth. Namely, if pixels are close in depth, then their influence is used in the upsampling. If pixels are not close in depth, then their influence may not be used in the upsampling. As mentioned above, in the illustrated example, the upsampling is going from 2 x 2 (i.e. 4 pixels) to
25 4 x 4 (i.e. 16 pixels). It should be understood, however, that the methods and techniques described herein may be applied to upsample any size small buffer to any size full size buffer. That is, the upsample ratio may vary. In some embodiments, the small white buffer may be chosen to have an aspect ratio similar to the full size white buffer. In general, the larger the small white buffer, the more detail that can be
30 obtained. However, increasing the size of the small white buffer generally increases costs.

In some embodiments, the upsampling (step 1230) may comprise bilateral

upsampling, or a bilateral upsample process. For example, in some embodiments, the process of zooming up the deferred buffer may involve bilateral upsampling techniques. In some embodiments, some of the above-described techniques for upsampling may be referred to as a bilateral upsample process. In some
5 embodiments, the techniques may include techniques used in bilateral filtering techniques. A bilateral filtering technique is discussed in the paper entitled "Image-Based Proxy Accumulation for Real-Time Soft Global Illumination", by Peter-Pike Sloan, Naga K. Govindaraju, Derek Nowrouzezahrai, John Snyder, Pacific Graphics 2007, October 2007, the entire contents and disclosure of which is hereby
10 incorporated by reference herein in its entirety.

As has been described, in some embodiments of the present invention there is provided a series of methods for rendering shadows, and for rendering image based projected shadows with multiple depth aware blurs, a purpose of which is to reduce aliasing artifacts and lend a soft appearance to shadows rendered into a multisampled
15 anti-aliased buffer. In some embodiments, any of the methods and techniques described above may be used for rendering shadows in real time. In some embodiments, aspects involve the use of a 3D caster and receiver grids, 2D grids, stenciling, cascades, categories, tiled shadow maps, etc. In some embodiments, aspects involve the use of a small white buffer or a deferred buffer, a blurring
20 technique that uses depth as a metric, and an upsample technique that takes the deferred buffer and zooms it up to a large buffer using depth as a metric. In some embodiments, the blurring comprises using depth values to determine whether or not a sample should be used in the blurring. In some embodiments, the upsampling
25 comprises using depth values to determine whether or not a sample should be used in the upsampling.

In some embodiments, the above-described methods and techniques may utilize or be combined with various shadow mapping and/or image based blurring algorithms. For example, in some embodiments, the above-described methods and techniques may include, or be combined with, or involve a synthesis of particular
30 sequences and applications of shadow mapping and image based depth-aware blurring algorithms.

The methods and techniques described herein may be utilized, implemented

and/or run on many different types of systems, including for example computers, game consoles, entertainment systems, etc. Referring to FIG. 14, there is illustrated an apparatus or system 1400 that may be used for any such implementations. One or more components of the apparatus or system 1400 may be used for implementing any system or device mentioned above, such as for example any of the above-mentioned computers, game consoles, entertainment systems, etc. However, the use of the apparatus or system 1400 or any portion thereof is certainly not required.

By way of example, the apparatus or system 1400 may include, but is not required to include, a central processing unit (CPU) 1402, a graphics processing unit (GPU) 1404, a random access memory (RAM) 1408, and a mass storage unit 1410, such as a disk drive or other type of memory. The apparatus or system 1400 may be coupled to, or integrated with, any of the other components described herein, such as a display 1412. The apparatus or system 1400 comprises an example of a processor based apparatus or system. The CPU 1402 and/or GPU 1404 may be used to execute or assist in executing the steps of the methods and techniques described herein, and various program content, images, shadows, lighting, etc. may be rendered on the display 1412.

The mass storage unit 1410 may include or comprise any type of computer readable storage or recording medium or media. The computer readable storage or recording medium or media may be fixed in the mass storage unit 1410, or the mass storage unit 1410 may optionally include removable storage media 1414, such as a digital video disk (DVD), Blu-ray disc, compact disk (CD), USB storage device, floppy disk, or other media. By way of example, the mass storage unit 1410 may comprise a disk drive, a hard disk drive, flash memory device, USB storage device, Blu-ray disc drive, DVD drive, CD drive, floppy disk drive, etc. The mass storage unit 1410 or removable storage media 1414 may be used for storing program code or macros that implement the methods and techniques described herein.

Thus, removable storage media 1414 may optionally be used with the mass storage unit 1410, which may be used for storing program code that implements the methods and techniques described herein. However, any of the storage devices, such as the RAM 1408 or mass storage unit 1410, may be used for storing such program code. For example, any of such storage devices may serve as a tangible computer

readable storage medium for storing or embodying a computer program for causing a console, system, computer, or other processor based system to execute or perform the steps of any of the methods, code, and/or techniques described herein. Furthermore, any of the storage devices, such as the RAM 1408 or mass storage unit 1410, may be used for storing any needed database(s).

In some embodiments, one or more of the embodiments, methods, approaches, and/or techniques described above may be implemented in a computer program executable by a processor based system. By way of example, such processor based system may comprise the processor based system 1400, or a computer, entertainment system, game console, graphics workstation, etc. Such computer program may be used for executing various steps and/or features of the above-described methods and/or techniques. That is, the computer program may be adapted to cause or configure a processor based system to execute and achieve the functions described above. For example, such computer program may be used for implementing any embodiment of the above-described steps or techniques for rendering shadows and/or rendering image based projected shadows with multiple depth aware blurs. As another example, such computer program may be used for implementing any type of tool or similar utility that uses any one or more of the above described embodiments, methods, approaches, and/or techniques. In some embodiments, the computer program may comprise a video game, role-playing game (RPG), other computer simulation, or system software such as an operating system, BIOS, macro, or other utility. In some embodiments, program code macros, modules, loops, subroutines, etc., within the computer program may be used for executing various steps and/or features of the above-described methods and/or techniques. In some embodiments, the computer program may be stored or embodied on a computer readable storage or recording medium or media, such as any of the computer readable storage or recording medium or media described herein.

In some embodiments, a processor-based apparatus may be used for executing or performing any of the above-described steps, methods, and/or techniques. For example, in some embodiments, the constructing the caster and receiver 3D grids, and/or the projecting and growing steps, may be performed by a processor-based apparatus.

Therefore, in some embodiments the present invention provides a computer program product comprising a medium for embodying a computer program for input to a computer and a computer program embodied in the medium for causing the computer to perform or execute steps comprising any one or more of the steps
5 involved in any one or more of the embodiments, methods, approaches, and/or techniques described herein.

For example, in some embodiments the present invention provides a computer readable storage medium storing one or more computer programs adapted to cause a processor based system to execute steps comprising identifying a set of shadow
10 casters and a set of shadow receivers in a frame of a scene that is lighted by at least a first light source, constructing a first three-dimensional (3D) grid, from the first light source's point of view, that fully encloses a first volume of the scene, comprising one or more shadow casters of the set of shadow casters that fall within the first volume of the scene, projecting each of the one or more shadow casters to generate one or more
15 projected shadow casters, growing each cell of the first 3D grid that includes at least a portion of the one or more projected casters, constructing a second three-dimensional (3D) grid, from the first light source's point of view, that fully encloses the first volume of the scene, comprising one or more shadow receivers of the set of shadow receivers that fall within the first volume of the scene, growing each cell of the second
20 3D grid that includes at least a portion of at least one of the one or more shadow receivers and for which a corresponding cell of the first 3D grid has been grown, marking for rendering each shadow receiver for which a cell of the second 3D grid has been grown, forming a third 3D grid by intersecting the first 3D grid with the second 3D grid such that the third 3D grid comprises the grown cells of the first grid and the grown cells of the second grid that intersect and marking for rendering into a
25 shadow map each shadow caster that intersects the third 3D grid.

In another embodiment, there is provided a computer readable storage medium storing one or more computer programs adapted to cause a processor based system to execute steps comprising: identifying one or more shadow casters and
30 corresponding shadow receivers in a frame of a scene that is lighted by at least a first light source; constructing a first three-dimensional (3D) grid, from the first light source's point of view, that fully encloses a first volume of the scene; projecting each

shadow caster that falls within the first 3D grid; growing each cell of the first 3D grid that includes at least a portion of a projected shadow caster; constructing a second three-dimensional (3D) grid, from the first light source's point of view, that fully encloses the first volume of the scene; growing each cell of the second 3D grid that
5 includes at least a portion of a shadow receiver and for which a corresponding cell of the first 3D grid has been grown; forming a third 3D grid by intersecting the first 3D grid with the second 3D grid; and marking for rendering into a shadow map each shadow caster that intersects the third 3D grid.

As another example, in some embodiments the present invention provides a
10 computer readable storage medium storing one or more computer programs adapted to cause a processor based system to execute steps comprising: identifying one or more groups of objects in a frame of content, wherein each group of objects includes one or more shadow casters and corresponding shadow receivers; for visible shadow receivers in the identified groups of objects, establishing depth information in a small
15 white buffer; for each of the identified groups of objects, rendering the shadow casters into a single depth map; blurring the small white buffer; and upsampling the small white buffer into a full sized white buffer, wherein the upsampling comprises using depth information as a metric.

The methods and embodiments described throughout this application are
20 described with respect to and may be used in 3D computer graphics for rendering shadows. In some embodiments, the method may be used in 3D computer graphics for rendering image based projected shadows. In some embodiments, the method may be used for rendering shadows in real time. However, it should be understood that the above methods and embodiments may be applied to any type of content, such
25 as for example video games, both online and offline game types, simulations, other types of videos, etc. In some embodiments, the above method may be used in any areas where shadowing may be desired and/or of importance.

While the invention herein disclosed has been described by means of specific embodiments, examples and applications thereof, numerous modifications and
30 variations could be made thereto by those skilled in the art without departing from the scope of the invention set forth in the claims.

CLAIMS

What is claimed is:

1. A computer readable storage medium storing one or more computer programs adapted to cause a processor based system to execute steps comprising:
 - identifying a set of shadow casters and a set of shadow receivers in a frame of a scene that is lighted by at least a first light source;
 - constructing a first three-dimensional (3D) grid, from the first light source's point of view, that fully encloses a first volume of the scene, comprising one or more shadow casters of the set of shadow casters that fall within the first volume of the scene;
 - projecting each of the one or more shadow casters to generate one or more projected shadow casters;
 - growing each cell of the first 3D grid that includes at least a portion of the one or more projected casters;
 - constructing a second three-dimensional (3D) grid, from the first light source's point of view, that fully encloses the first volume of the scene, comprising one or more shadow receivers of the set of shadow receivers that fall within the first volume of the scene;
 - growing each cell of the second 3D grid that includes at least a portion of at least one of the one or more shadow receivers and for which a corresponding cell of the first 3D grid has been grown;
 - marking for rendering each shadow receiver for which a cell of the second 3D grid has been grown;
 - forming a third 3D grid by intersecting the first 3D grid with the second 3D grid such that the third 3D grid comprises the grown cells of the first grid and the grown cells of the second grid that intersect; and
 - marking for rendering into a shadow map each shadow caster that intersects the third 3D grid.

2. The computer readable storage medium of claim 1, wherein the one or more computer programs are further adapted to cause the processor based system to execute steps comprising:

- partitioning a screen into a two-dimensional (2D) grid;
- projecting each cell of the 2D grid from a near plane to a far plane to form a corresponding frustum;
- for each cell of the 2D grid, determining whether or not any grown cells of the third 3D grid fall within the corresponding frustum; and
- excluding each cell of the 2D grid from receiving shadows for the first light source for which no grown cells of the third 3D grid fall within the corresponding frustum.

3. The computer readable storage medium of claim 2, wherein the one or more computer programs are further adapted to cause the processor based system to execute steps comprising:

- for each cell of the 2D grid for which one or more grown cells of the third 3D grid do fall within the corresponding frustum, determining an identity and a volume of the one or more grown cells of the third 3D grid that fall within the corresponding frustum; and

- for each cell of the 2D grid for which one or more grown cells of the third 3D grid do fall within the corresponding frustum, excluding screen pixels whose z-buffer depths fall outside the determined volume of the one or more grown cells of the third 3D grid that fall within the corresponding frustum.

4. The computer readable storage medium of claim 2 or 3, wherein the one or more computer programs are further adapted to cause the processor based system to execute steps comprising

- determining a 3D AABB intersection volume for each cell of the 2D grid for which one or more grown cells of the third 3D grid fall within the corresponding frustum.

5. The computer readable storage medium of any one of claims 1 to 3, wherein the one or more computer programs are further adapted to cause the processor based system to execute steps comprising:

dividing the scene into a plurality of cascades that are broken up along depth from a rendering camera's point of view;

wherein the first volume of the scene comprises one of the plurality of cascades.

6. The computer readable storage medium of claim 1, wherein the one or more computer programs are further adapted to cause the processor based system to execute steps comprising:

rendering the marked shadow casters into the shadow map.

7. The computer readable storage medium of claim 6, wherein the one or more computer programs are further adapted to cause the processor based system to execute steps comprising:

tiling the shadow map.

8. The computer readable storage medium of any one of claims 1 to 3 or 6, wherein the one or more computer programs are further adapted to cause the processor based system to execute steps comprising:

performing a depth prepass on rendering geometry in the frame and storing depth information in a depth buffer.

9. The computer readable storage medium of claim 8, wherein the one or more computer programs are further adapted to cause the processor based system to execute steps comprising:

reconstructing shadow receivers by performing steps comprising sampling from the depth buffer and sampling from the shadow map.

10. The computer readable storage medium of claim 9, wherein the one or more computer programs are further adapted to cause the processor based system to

execute steps comprising:

during the depth prepass, for each of one or more of pixels, writing into a stencil a value indicating a type of receiver that the pixel is; and

discriminating which pixels receive shadow based on the corresponding values in the stencil.

11. The computer readable storage medium of claim 9, wherein the one or more computer programs are further adapted to cause the processor based system to execute steps comprising:

dividing the one or more shadow casters into a plurality of categories; and

discriminating which pixels receive shadow based on the category of a relevant shadow caster.

12. The computer readable storage medium of claim 9, wherein the one or more computer programs are further adapted to cause the processor based system to execute steps comprising:

writing resultant shadow values into a deferred buffer that was initially set to white.

13. The computer readable storage medium of claim 12, wherein the one or more computer programs are further adapted to cause the processor based system to execute steps comprising:

dilating the deferred buffer.

14. The computer readable storage medium of claim 12, wherein the one or more computer programs are further adapted to cause the processor based system to execute steps comprising:

rendering the marked shadow receivers by sampling the deferred buffer.

15. The computer readable storage medium of claim 14, wherein the one or more computer programs are further adapted to cause the processor based system to execute steps comprising:

displaying the rendered shadow receivers on a display.

16. The computer readable storage medium of any one of claims 1 to 3, wherein the one or more computer programs are further adapted to cause the processor based system to execute steps comprising:

collecting per draw call bounding information for the set of shadow caster and set of shadow receivers.

17. The computer readable storage medium of claim 16, wherein the one or more computer programs are further adapted to cause the processor based system to execute steps wherein the step of constructing the first 3D grid comprises placing the one or more shadow casters within the first 3D grid based on the per draw call bounding information.

18. The computer readable storage medium of claim 16, wherein the one or more computer programs are further adapted to cause the processor based system to execute steps wherein the step of constructing the second 3D grid comprises placing the one or more shadow receivers within the first 3D grid based on the per draw call bounding information.

19. A method, comprising:

identifying a set of shadow casters and a set of shadow receivers in a frame of a scene that is lighted by at least a first light source;

constructing a first three-dimensional (3D) grid, from the first light source's point of view, that fully encloses a first volume of the scene, comprising one or more shadow casters of the set of shadow casters that fall within the first volume of the scene;

projecting each of the one or more shadow casters to generate one or more projected shadow casters;

growing each cell of the first 3D grid that includes at least a portion of the one or more projected casters;

constructing a second three-dimensional (3D) grid, from the first light source's

point of view, that fully encloses the first volume of the scene, comprising one or more shadow receivers of the set of shadow receivers that fall within the first volume of the scene;

growing each cell of the second 3D grid that includes at least a portion of at least one of the one or more shadow receivers and for which a corresponding cell of the first 3D grid has been grown;

marking for rendering each shadow receiver for which a cell of the second 3D grid has been grown;

forming a third 3D grid by intersecting the first 3D grid with the second 3D grid such that the third 3D grid comprises the grown cells of the first grid and the grown cells of the second grid that intersect; and

marking for rendering into a shadow map each shadow caster that intersects the third 3D grid.

20. The method of claim 19, further comprising:

partitioning a screen into a two-dimensional (2D) grid;

projecting each cell of the 2D grid from a near plane to a far plane to form a corresponding frustum;

for each cell of the 2D grid, determining whether or not any grown cells of the third 3D grid fall within the corresponding frustum; and

excluding each cell of the 2D grid from receiving shadows for the first light source for which no grown cells of the third 3D grid fall within the corresponding frustum.

21. The method of claim 20, further comprising:

for each cell of the 2D grid for which one or more grown cells of the third 3D grid do fall within the corresponding frustum, determining an identity and a volume of the one or more grown cells of the third 3D grid that fall within the corresponding frustum; and

for each cell of the 2D grid for which one or more grown cells of the third 3D grid do fall within the corresponding frustum, excluding screen pixels whose z-buffer depths fall outside the determined volume of the one or more grown cells of the third

3D grid that fall within the corresponding frustum.

22. An apparatus, comprising:

a display; and

a processor based system coupled to the display, wherein the processor based system is configured to execute steps comprising:

identifying a set of shadow casters and a set of shadow receivers in a frame of a scene that is lighted by at least a first light source;

constructing a first three-dimensional (3D) grid, from the first light source's point of view, that fully encloses a first volume of the scene, comprising one or more shadow casters of the set of shadow casters that fall within the first volume of the scene;

projecting each of the one or more shadow casters to generate one or more projected shadow casters;

growing each cell of the first 3D grid that includes at least a portion of the one or more projected casters;

constructing a second three-dimensional (3D) grid, from the first light source's point of view, that fully encloses the first volume of the scene, comprising one or more shadow receivers of the set of shadow receivers that fall within the first volume of the scene;

growing each cell of the second 3D grid that includes at least a portion of at least one of the one or more shadow receivers and for which a corresponding cell of the first 3D grid has been grown;

marking for rendering each shadow receiver for which a cell of the second 3D grid has been grown;

forming a third 3D grid by intersecting the first 3D grid with the second 3D grid such that the third 3D grid comprises the grown cells of the first grid and the grown cells of the second grid that intersect; and

marking for rendering into a shadow map each shadow caster that intersects the third 3D grid.

23. The apparatus of claim 22, wherein the processor based system is further

configured to execute steps comprising:

- partitioning a screen into a two-dimensional (2D) grid;
- projecting each cell of the 2D grid from a near plane to a far plane to form a corresponding frustum;
- for each cell of the 2D grid, determining whether or not any grown cells of the third 3D grid fall within the corresponding frustum; and
- excluding each cell of the 2D grid from receiving shadows for the first light source for which no grown cells of the third 3D grid fall within the corresponding frustum.

24. The apparatus of claim 23, wherein the processor based system is further configured to execute steps comprising:

- for each cell of the 2D grid for which one or more grown cells of the third 3D grid do fall within the corresponding frustum, determining an identity and a volume of the one or more grown cells of the third 3D grid that fall within the corresponding frustum; and
- for each cell of the 2D grid for which one or more grown cells of the third 3D grid do fall within the corresponding frustum, excluding screen pixels whose z-buffer depths fall outside the determined volume of the one or more grown cells of the third 3D grid that fall within the corresponding frustum.

25. A computer readable storage medium storing one or more computer programs adapted to cause a processor based system to execute steps comprising:

- identifying one or more groups of objects in a frame of content, wherein each group of objects includes one or more shadow casters and corresponding shadow receivers;

- for visible shadow receivers in the identified groups of objects, establishing depth information in a small white buffer;

- for each of the identified groups of objects, rendering the shadow casters into a single depth map;

- blurring the small white buffer; and

- upsampling the small white buffer into a full sized white buffer, wherein the

upsampling comprises using depth information as a metric.

26. The computer readable storage medium of claim 25, wherein the one or more computer programs are further adapted to cause the processor based system to execute steps comprising:

rendering the shadow receivers in the identified groups of objects by sampling the upsampled full sized white buffer.

27. The computer readable storage medium of claim 25, wherein the one or more computer programs are further adapted to cause the processor based system to execute steps comprising:

performing a first depth prepass with all opaque and cutout geometry in the content;

wherein the upsampling is performed using high resolution depth information stored in the first depth prepass to weight each of a number of bilinear low resolution samples.

28. The computer readable storage medium of claim 25, wherein the upsampling comprises setting weights to zero for each sample having a depth that is beyond a certain distance from a high resolution sample point's depth.

29. The computer readable storage medium of claim 25, wherein the upsampling comprises weighting based on distance each sample having a depth that is within a certain distance from a high resolution sample point's depth.

30. The computer readable storage medium of any one of claims 25 to 29, wherein the blurring comprises performing blurring with a depth aware low pass filter.

31. The computer readable storage medium of any one of claims 25 to 29, wherein the blurring comprises not including an influence of samples having a depth that is beyond a certain distance from a sample point.

32. The computer readable storage medium of any one of claims 25 to 29, wherein the blurring comprises weighting based on distance each sample having a depth that is within a certain distance from a sample point.

33. The computer readable storage medium of any one of claims 25 to 29, wherein the establishing depth information in a small white buffer comprises rendering a depth prepass to the small white buffer for visible shadow receivers in the identified groups of objects.

34. The computer readable storage medium of any one of claims 25 to 29, wherein the one or more computer programs are further adapted to cause the processor based system to execute steps comprising:

for each of the identified groups of objects, rendering the shadow receivers to the small white buffer;

wherein the rendering the shadow receivers to the small white buffer comprises using a certain amount of filtering samples.

35. The computer readable storage medium of claim 25, wherein the upsampling comprises using depth values to determine whether or not a sample should be used in the upsampling.

36. The computer readable storage medium of any one of claims 25 to 29 or 35, wherein the blurring comprises using depth values to determine whether or not a sample should be used in the blurring.

37. A method, comprising:

identifying one or more groups of objects in a frame of content, wherein each group of objects includes one or more shadow casters and corresponding shadow receivers;

for visible shadow receivers in the identified groups of objects, establishing depth information in a small white buffer;

for each of the identified groups of objects, rendering the shadow casters into a single depth map;

blurring the small white buffer;

upsampling the small white buffer into a full sized white buffer, wherein the upsampling comprises using depth information as a metric and is performed by a processor-based apparatus.

38. The method of claim 37, further comprising:

rendering the shadow receivers in the identified groups of objects by sampling the upsampled full sized white buffer; and

displaying the rendered shadow receivers on a display.

39. The method of claim 37, further comprising:

performing a first depth prepass with all opaque and cutout geometry in the content;

wherein the upsampling is performed using high resolution depth information stored in the first depth prepass to weight each of a number of bilinear low resolution samples.

40. The method of claim 37, wherein the upsampling comprises setting weights to zero for each sample having a depth that is beyond a certain distance from a high resolution sample point's depth.

41. The method of claim 37, wherein the upsampling comprises weighting based on distance each sample having a depth that is within a certain distance from a high resolution sample point's depth.

42. The method of any one of claims 37 to 41, wherein the blurring comprises performing blurring with a depth aware low pass filter.

43. The method of any one of claims 37 to 41, wherein the blurring comprises not including an influence of samples having a depth that is beyond a certain distance

from a sample point.

44. The method of any one of claims 37 to 41, wherein the blurring comprises weighting based on distance each sample having a depth that is within a certain distance from a sample point.

45. The method of any one of claims 37 to 41, wherein the establishing depth information in a small white buffer comprises rendering a depth prepass to the small white buffer for visible shadow receivers in the identified groups of objects.

46. The method of any one of claims 37 to 41, further comprising:
for each of the identified groups of objects, rendering the shadow receivers to the small white buffer;
wherein the rendering the shadow receivers to the small white buffer comprises using a certain amount of filtering samples.

47. The method of claim 37, wherein the upsampling comprises using depth values to determine whether or not a sample should be used in the upsampling.

48. The method of any one of claims 37 to 41 or 47, wherein the blurring comprises using depth values to determine whether or not a sample should be used in the blurring.

49. An apparatus, comprising:
a display; and
a processor based system coupled to the display, wherein the processor based system is configured to execute steps comprising:
identifying one or more groups of objects in a frame of content, wherein each group of objects includes one or more shadow casters and corresponding shadow receivers;
for visible shadow receivers in the identified groups of objects,
establishing depth information in a small white buffer;

for each of the identified groups of objects, rendering the shadow casters into a single depth map;

blurring the small white buffer; and

upsampling the small white buffer into a full sized white buffer, wherein the upsampling comprises using depth information as a metric.

50. The apparatus of claim 49, wherein the processor based system is further configured to execute steps comprising:

rendering the shadow receivers in the identified groups of objects by sampling the upsampled full sized white buffer; and

displaying the rendered shadow receivers on the display.

51. The apparatus of claim 49, wherein the upsampling comprises setting weights to zero for each sample having a depth that is beyond a certain distance from a high resolution sample point's depth.

52. The apparatus of claim 49, wherein the upsampling comprises weighting based on distance each sample having a depth that is within a certain distance from a high resolution sample point's depth.

53. The apparatus of any one of claims 49 to 52, wherein the blurring comprises performing blurring with a depth aware low pass filter.

54. The apparatus of any one of claims 49 to 52, wherein the blurring comprises not including an influence of samples having a depth that is beyond a certain distance from a sample point.

55. The apparatus of any one of claims 49 to 52, wherein the blurring comprises weighting based on distance each sample having a depth that is within a certain distance from a sample point.

56. The apparatus of any one of claims 49 to 52, wherein the establishing

depth information in a small white buffer comprises rendering a depth prepass to the small white buffer for visible shadow receivers in the identified groups of objects.

57. The apparatus of claim 49, wherein the upsampling comprises using depth values to determine whether or not a sample should be used in the upsampling.

58. The apparatus of any one of claims 49 to 52 or 57, wherein the blurring comprises using depth values to determine whether or not a sample should be used in the blurring.




FIG. 1




FIG. 2




FIG. 3




FIG. 4




FIG. 5




FIG. 6




FIG. 7

800 ↗




FIG. 8




FIG. 9




FIG. 10




FIG. 11




FIG. 12




FIG. 13




FIG. 14

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2010/035676

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - G06T 15/60 (2010.01)

USPC - 345/426

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC(8) - G06T 11/40, 15/00, 15/60 (2010.01)

USPC - 345/421, 422, 426

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)

MicroPatent and Google Patents

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	BOULANGER. Real-Time Realistic Rendering of Nature Scenes with Dynamic Lighting. M.S. INRIA, University of Rennes I, France, 2005. [Retrieved on 2010-09-15]. Retrieved from the Internet: <URL:http://www.kevinboulanger.net/publications/thesisKevinBoulanger.pdf> pp 1-191.	1-58
Y	US 7,142,709 B2 (GIRARD) 28 November 2006 (28.11.2006) entire document	1-24
Y	US 5,347,374 A (FUSS et al) 13 September 1994 (13.09.1994) entire document	5
Y	US 2007/0002047 A1 (DESGRANGES et al) 04 January 2007 (04.01.2007) entire document	13
Y	US 2006/0221250 A1 (ROSSBACH et al) 05 October 2006 (05.10.2006) entire document	25-58

☐

Further documents are listed in the continuation of Box C.

☐

* 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

16 September 2010

Date of mailing of the international search report

23 SEP 2010

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
P.O. Box 1450, Alexandria, Virginia 22313-1450

Facsimile No. 571-273-3201

Authorized officer:

Blaine R. Copenheaver

PCT Helpdesk: 571-272-4300
PCT OSP: 571-272-7774

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2010/035676

Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:
2. ☐ Claims Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:
3. ☐ Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

See extra sheet.

1. ☒ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. ☐ As all searchable claims could be searched without effort justifying additional fees, this Authority did not invite payment of additional fees.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
4. ☐ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

Remark on Protest

- ☐ The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
- ☐ The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
- ☒ No protest accompanied the payment of additional search fees.

Continuation of Box III.

This application contains the following inventions or groups of inventions which are not so linked as to form a single general inventive concept under PCT Rule 13.1. In order for all inventions to be examined, the appropriate additional examination fees must be paid.

Group I, claims 1-24, drawn to a computer readable storage medium, method and apparatus comprising constructing a first three-dimensional (3D) grid, from a first light source's point of view, that fully encloses a first volume of the scene, comprising one or more shadow casters of the set of shadow casters that fall within the first volume of the scene; projecting each of the one or more shadow casters to generate one or more projected shadow casters; growing each cell of the first 3D grid that includes at least a portion of the one or more projected casters; constructing a second three-dimensional (3D) grid, from the first light source's point of view, that fully encloses the first volume of the scene, comprising one or more shadow receivers of the set of shadow receivers that fall within the first volume of the scene; growing each cell of the second 3D grid that includes at least a portion of at least one of the one or more shadow receivers and for which a corresponding cell of the first 3D grid has been grown; marking for rendering each shadow receiver for which a cell of the second 3D grid has been grown; forming a third 3D grid by intersecting the first 3D grid with the second 3D grid such that the third 3D grid comprises the grown cells of the first grid and the grown cells of the second grid that intersect; and marking for rendering into a shadow map each shadow caster that intersects the third 3D grid.

Group II, claims 25-58, drawn to a computer readable storage medium, method and apparatus comprising establishing depth information in a small white buffer for visible shadow receivers in the identified groups of objects; for each of the identified groups of objects, rendering the shadow casters into a single depth map; blurring the small white buffer; and upsampling the small white buffer into a full sized white buffer, wherein the upsampling comprises using depth information as a metric.

The inventions listed as Groups I-II do not relate to a single general inventive concept under PCT Rule 13.1 because, under PCT Rule 13.2, they lack the same or corresponding special technical features for the following reasons: the special technical feature of the Group I invention: as claimed therein is not present in the invention of Group II. The special technical feature of the Group II invention: as claimed therein is not present in the invention of Groups I.

Groups I and II lack unity of invention because even though the inventions of these groups require the technical feature of a processor based system to execute steps comprising identifying a set of shadow casters and a set of shadow receivers in a frame of a scene that is lighted by at least a first light source, this technical feature is not a special technical feature as it does not make a contribution over the prior art in view of US 7,142,709 B2 (GIRARD) 28 November 2006 (28.11.2006) column 6, lines 25-45; column 13, lines 40-45; column 14, lines 17-35; figure 5.

Since none of the special technical features of the Group I or II inventions are found in more than one of the inventions, unity of invention is lacking.