



- (51) International Patent Classification:
G06T 11/20 (2006.01)
- (21) International Application Number:
PCT/US2014/026799
- (22) International Filing Date:
13 March 2014 (13.03.2014)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
13/831,309 14 March 2013 (14.03.2013) US
- (71) Applicant: DREAMWORKS ANIMATION LLC
[US/US]; 1000 Flower Street, Glendale, CA 91201 (US).
- (72) Inventor: SMYTH, Evan, P.; c/o DREAMWORKS ANIMATION LLC, 1000 Flower Street, Glendale, CA 91201 (US).
- (74) Agents: YIM, Peter, J. et al.; Morrison & Foerster LLP, 425 Market Street, San Francisco, CA 94105-2482 (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, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK,

[Continued on next page]

(54) Title: CONSERVATIVE PARTITIONING FOR RENDERING A COMPUTER-GENERATED ANIMATION




FIG. 6

(57) Abstract: A system for partitioning a set of assets, where each asset represents a computer-generated object associated with a computer-generated scene. A dependency graph comprising a plurality of interconnected nodes including an organizer node is accessed. The set of assets identified by an input of a predicate test of the organizer node are accessed. It is determined if the at least one predicate test can be evaluated using the set of assets. If the at least one predicate test can be evaluated, one or more partition assets are identified and passed to a next node. If the at least one predicate test cannot be evaluated, a conservative set of assets is identified and passed to the next node, wherein the conservative set of assets is the same set of assets identified by the input of the predicate test.



EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU,
LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK,
SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ,
GW, KM, ML, MR, NE, SN, TD, TG).

Published:

— *without international search report and to be republished
upon receipt of that report (Rule 48.2(g))*

CONSERVATIVE PARTITIONING FOR RENDERING A COMPUTER-GENERATED ANIMATION

BACKGROUND

1. Field

[0001] The present disclosure relates to computer-generated animation and, more specifically, to rendering an image of a computer-generated scene using a dependency graph.

2. Related Art

[0002] A computer-generated animation image can be created by rendering one or more computer-generated objects to depict a scene in the image. Light sources, materials, textures, and other visual effects can be associated with the computer-generated objects to create a realistic visual appearance for the animation image. A dependency graph can be used to define the relationships between assets (which represent the computer-generated objects) and a variety of visual effects as part of the rendering process. The dependency graph typically includes one or more interconnected nodes associating the assets with one or more visual effects, where a node wire can pass the assets and visual effects from node-to-node for processing. The output of the dependency graph can be used to create a rendered image of the scene.

[0003] In a typical implementation, a dependency graph is used to render a scene using a subset of assets that is less than a complete set of assets associated with the scene. For example, some of the objects associated with the scene may be out of the camera's view for a particular rendered scene and, therefore, do not need to be rendered. Because loading assets into memory takes time and resources, it is generally advantageous to only load the subset of assets that are required to render the scene. However, using traditional techniques, the precise subset of assets that are necessary for rendering is not known until all the asset information has been loaded so that the dependency graph can be evaluated. Traditional techniques may result in the loading of assets (including geometry of the computer-generated objects) that are not used, and as a result, consume more computer resources than necessary. It is therefore desirable to create a system for evaluating dependency graphs with less than a complete set of assets associated with a scene.

SUMMARY

[0004] One exemplary embodiment includes a computer-implemented method for partitioning a set of assets. Each asset represents a computer-generated object associated with a computer-generated scene. A dependency graph is accessed. The dependency graph comprises a plurality of interconnected nodes, wherein one of the interconnected nodes includes an organizer node having at least one predicate test. The set of assets identified by an input of the predicate test of the organizer node are accessed. It is determined if the at least one predicate test can be evaluated using the set of assets. If the at least one predicate test can be evaluated, a logical operation is performed on the set of assets to identify one or more partition assets, wherein the one or more partition assets are a subset of the set of assets. The one or more partition assets identified by an output of the organizer node are then passed to a next node in the dependency graph. If the at least one predicate test cannot be evaluated, a conservative set of assets identified by the output of the organizer node is passed to the next node, wherein the conservative set of assets is the same set of assets identified by the input of the predicate test.

[0005] In some embodiments, the set of assets includes asset geometry and asset attributes for one or more assets of the set of assets. In some cases, the set of assets does not include asset geometry for one or more assets of the set of assets.

[0006] In some cases, the at least one predicate test can be evaluated if a first asset attribute is required for the logical operation and the first asset attribute is able to be accessed at the input to the predicate test. In some cases, the at least one predicate test cannot be evaluated if a first asset attribute is required for the logical operation and the first asset attribute is not able to be accessed at the input to the predicate test. In some cases, the at least one predicate test cannot be evaluated if the asset geometry is required for the logical operation and the asset geometry is able to be accessed at the input to the predicate test. In some cases, the at least one predicate test cannot be evaluated because at least some information for one or more assets of the set of assets is not loaded in computer memory.

[0007] In some embodiments, one or more nodes of the plurality of nodes assigns a visual effect to the set of assets. The visual effect may be a virtual light source.

[0008] In one exemplary embodiment, the one or more partition assets are passed to a render node of the dependency graph. A rendered image of the computer-generated objects represented by the one or more partition assets is computed using the render node.

[0009] One exemplary embodiment includes a computer-implemented method for partitioning a set of assets, wherein each asset represents a computer-generated object associated with a computer-generated scene. A dependency graph is accessed. The dependency graph comprises a plurality of interconnected nodes, including a render node associated with a rendering pass. A render node of the dependency graph is identified, the render node associated with a rendering pass of the computer-generated scene. A set of nodes of the dependency graph that are connected to the render node and are upstream from the render node are identified. A first pass through the dependency graph is performed through the set of nodes to identify a set of partition assets. Asset information is loaded for the set of partition assets. A second pass through the dependency graph through the set of nodes using the asset information for the set of partition assets is performed. A rendered image is computed based on the second pass through the dependency graph.

BRIEF DESCRIPTION OF THE FIGURES

[0010] The present application can be best understood by reference to the following description taken in conjunction with the accompanying drawing figures, in which like parts may be referred to by like numerals.

[0011] FIG. 1 depicts a set of computer-generated objects associated with a scene.

[0012] FIG. 2 depicts a set of computer-generated objects associated with a scene with lighting effects applied.

[0013] FIG. 3 depicts a rendered image of a scene.

[0014] FIG. 4 depicts a schematic representation of a render setup graph.

[0015] FIG. 5 depicts a schematic representation of an organizer node of a render setup graph.

[0016] FIG. 6 depicts an exemplary process for evaluating a dependency graph.

[0017] FIG. 7 depicts an exemplary process for rendering an image using a dependency graph.

[0018] FIG. 8 depicts an exemplary workstation computer system.

DETAILED DESCRIPTION

[0019] The following description is presented to enable a person of ordinary skill in the art to make and use the various embodiments. Descriptions of specific devices, techniques, and applications are provided only as examples. Various modifications to the examples described herein will be readily apparent to those of ordinary skill in the art, and the general principles defined herein may be applied to other examples and applications without departing from the spirit and scope of the present technology. Thus, the disclosed technology is not intended to be limited to the examples described herein and shown, but is to be accorded the scope consistent with the claims.

[0020] A computer-generated animation is typically created by rendering a sequence of images, each image depicting a computer-generated scene composed of one or more computer-generated objects. Light sources, materials, textures, and other visual effects are associated with the computer-generated objects to create a realistic visual appearance for the computer-generated scene.

[0021] FIG. 1 depicts an exemplary computer-generated scene having multiple computer-generated objects positioned in a three-dimensional space. For purposes of this discussion, a computer-generated scene generally refers to the three-dimensional space that can be filmed using a virtual camera, and may also be referred to generically as a scene. As shown in FIG. 1, the scene 100 includes surface models of a teapot 151, a sphere 152, a cylinder 153, a cone 153, and a cube 154. In a typical computer animation, the computer-generated objects include one or more computer-animated characters that have been posed or manipulated for purposes of generating a computer-animated film sequence. In this example, the computer-generated objects are surface models defined using surface geometry. In other examples, the computer-generated objects may be defined using solid geometry or defined using other three-dimensional modeling techniques.

[0022] As shown in FIG. 1, a virtual camera 104 is positioned to view a portion of the scene 100. Typically, the virtual camera 104 has a field of view that is defined by a camera

frustum that projects away from the lens of the virtual camera 104. In general, computer-generated objects that are outside of the camera frustum are not directly seen by the camera. The positions of the virtual camera 104 and computer-generated objects 151-154 in the scene 100 are typically determined by a human operator, such as an animator or director.

[0023] As described in more detail below, lights and other visual effects may also be applied to the scene 100 to give it a more realistic appearance. FIG. 2 depicts the scene 100 with lighting effects applied to the computer-generated objects 151-154. As shown in FIG. 2, the computer-generated objects have been illuminated by a diffuse light source so that the computer-generated objects 151-154 are illuminated from the camera-side. Each computer-generated object 151-154 is shaded in accordance with the diffuse light source and casts a shadow on the floor and walls of the scene 100. The scene 100 is also illuminated by one or more ambient light sources to provide an overall illumination lighting effect to the scene 100. Other light sources that can be used to illuminate the scene 100 include, for example, point lights, spot lights, and area lights.

[0024] FIG. 3 depicts a rendered image 101 of the portion of the scene 100 in view of the virtual camera 104. As shown in FIG. 3, the image 101 depicts a subset of the computer-generated objects, specifically the teapot 151, cone 153, and cube 154. The computer-generated objects 151, 153, and 154 in FIG. 3 have been rendered using multiple light sources and one or more types of surface shaders for each computer-generated object to produce a realistic-looking image of the scene 100. Surface shaders are used to simulate the optical properties of the surface of the computer-generated objects and define the color, texture, specular properties, and other optical characteristics of the surface of the computer-generated objects. A surface shader may use, for example, a bidirectional reflectance distribution function (BRDF) or other technique for simulating the reflection of light incident on a computer-generated surface. The surface shaders may also be used to define a group of optical properties to simulate a material (e.g., fur, skin, or metal). Typically, the configuration of the light sources, surface shaders, and surface materials are configurable for each computer-generated scene and may be referred to generally as a lighting configuration.

[0025] The lighting configuration may be designated on an object-by-object basis or by set of objects identified by a shared asset attribute. For example, a spot-light light source may be associated only with the main characters in the foreground of the scene and may not illuminate other objects. Alternatively, an ambient light might be associated with only the

plants in the scene and have no effect on the main characters. This approach simplifies the computations required to render a scene and also provides the animator or user with more flexibility when configuring the visual effects in a scene. However, the additional flexibility can make it difficult to manage a large number of computer-generated objects and their associated visual effects.

[0026] As previously mentioned, a dependency graph can be used to manage the associations between computer-generated objects and visual effects. For purposes of this discussion, a particular type of dependency graph, called a render setup graph, is used to manage the associations between computer-generated objects and visual effects used to render the scene. FIG. 4 depicts an exemplary visualization of a render setup graph 200. The render setup graph 200 may be displayed on a computer display and manipulated using a graphical user interface and computer I/O hardware, as described in more detail in Section 3, below. The render setup graph 200 is generally configurable by an animator or user and can be used to create multiple lighting scenarios for a scene.

[0027] The render setup graph 200 typically accepts as input, a set of computer-generated objects represented by assets located in asset storage 202. The asset storage 202 may be the computer memory of the workstation computer system that is managing the render setup graph 200. The asset storage 202 may also be connected to a remote, centralized computer-storage database that can be accessed by multiple users. A centralized database may improve overall storage efficiency and helps ensure that the computer-generated animation is using a common version of the computer-generated objects.

[0028] In the render setup graph 200, visual effects, such as light sources, shaders, and materials are assigned to the assets at various nodes 204A-E. In this example, the render setup graph 200 is evaluated from left to right with each node 204A-E assigning one or more visual effects to a set of assets, each asset representing a computer-generated object in the scene. Node wires 206 are used to pass assets and associated visual effects between elements of the render setup graph 200. Eventually, the assets and their associated visual effects are passed to one of the render nodes 208A-C. In this example, each render node 208A-C contains the completed lighting configuration and can be stored or used to create a rendered image of the scene.

[0029] Render setup graph 200 also includes a special type of node called an organizer node (210A-B), which is used to select an asset or subset of assets for downstream node operations. The selection of assets for downstream nodes, also called partitioning, is useful for assigning visual effects based on attributes of the assets. For example, one of the organizer nodes 210A-B may be used to select or identify all of the assets in the scene that are within the view angle or frustum of the camera. The organizer nodes 210A-B may also be used to select or identify assets based on asset name, location, size, or other asset attribute. The subset of assets that are identified by the organizer nodes 210A-B are also referred to herein as a partition or subset of partition assets.

[0030] Typically, the organizer nodes 210A-B have multiple outputs, each output representing a subset of partition assets that can be passed to another element in the render setup graph 200. In a typical implementation, the subset of partition assets at an output (e.g., 212A) of one of the organizer node 210A-B is a reduced set of assets as compared to the set of assets received at the input 211 of the organizer node 210A-B. In this way, the assets are culled or reduced in number as they pass through the organizer nodes 210A-B, to the respective downstream nodes 204C-E. Eventually, the subset of partition assets is passed to the render nodes 208A-C and used to compute a rendered image of the scene.

[0031] Render setup graphs typically include other elements that provide a variety of functions and operations on the assets as they propagate through the elements. For example, a render setup graph may also include splitter, combiner, and other routing elements. These other elements are omitted from this discussion and the example depicted in FIG. 4 for clarity. The techniques discussed below also apply to render setup graphs having additional elements or fewer elements to those depicted in the render setup graph 200 of FIG. 4.

[0032] In general, a render setup graph can be used to define multiple rendering passes for a scene. Each rendering pass may combine a subset of partition assets with a unique set of visual effects. For example, a rendering pass may be configured to render only the assets in the background of the scene. A rendering pass may also be configured to render a subset of partition assets using a simplified set of light sources and shaders to reduce the rendering computations that are necessary to compute the rendered image of the scene.

[0033] The render setup graph 200 depicted in FIG. 4 depicts three exemplary rendering passes, as indicated by the presence of three render nodes 208A-C. The lighting

configurations for the subsets of partition assets of rendering pass are defined by the respective nodes 204A-C and organizer nodes 210A-B that are connected to each respective render node 208A-C by the node wires 206.

[0034] The ability to define and configure rendering passes allows the animator or user to develop different lighting configurations without having to work with the entire set of assets that are associated with a scene. Multiple rendering passes also allow the animator or user to work with a simplified version of the scene that may omit resource-intensive visual effects that are used in the final production rendering pass. In some cases, multiple rendering passes may be combined to produce a final rendered image of the complete scene.

[0035] In general, because the assets are selected or partitioned by the organizer nodes 210A-B, any single rendering pass (corresponding to one of the render nodes 208A-C) does not use all of assets that are associated with a scene. Because the assets are typically stored remotely in a centralized computer-storage database, significant computer resources are consumed retrieving and creating local copies of assets. As previously mentioned, by not loading all of the assets for every rendering pass, a savings in computer resources can be realized.

[0036] However, the precise set of assets that are necessary for a rendering pass may not be known until the render setup graph 200 is evaluated. In particular, organizer nodes 210A-B may not be able to select or partition the assets for downstream elements if all of the asset information is not loaded and provided as input to the render setup graph 200. However, as described in more detail below, organizer nodes 110A-B can be specially configured to output a partially culled, but conservative set of assets to downstream elements of the render setup graph 200. In some cases, the conservative set of asserts is over-inclusive of the precise set of assets that is required for the rendering pass, but is reduced with respect to the entire set of assets associated with a scene. Accordingly, the system and techniques described can be used to evaluate a rendering pass without loading all of the assets associated the scene being rendered.

1. Evaluating Organizer Nodes Using Incomplete Asset Information

[0037] FIG. 6 depicts an exemplary process 1000 for evaluating a dependency graph, such as a render setup graph. Process 1000 may be performed using a workstation computer system, as described below with respect to FIG. 8 as part of a computer-animation process.

Exemplary process 1000 may be performed, for example, after the computer-generated objects (e.g., animated characters) have been placed or posed in the scene and stored as assets in a centralized storage database. Process 1000 is typically performed before an image of the scene has been rendered.

[0038] Process 1000 may provide advantages over other techniques by requiring fewer computing resources, particularly in scenarios where less than all of the asset information has been loaded in the computer memory of the workstation computer. As described above with respect to FIG. 4, the assets used by the render setup graph 200 are stored in asset storage 202, which may retrieve information from a centralized computer-storage database that can be accessed by multiple users over a computer network. Loading assets from the database into workstation memory over the network consumes a nontrivial amount computer resources. Thus, when using a centralized storage configuration, loading fewer assets into workstation memory may reduce network traffic, reduce loads on the centralized database, and may reduce asset loading time.

[0039] Process 1000 may be used when computing a rendering pass of a computer-generated scene for an animated film sequence using a dependency graph. For purposes of the following discussion, examples of process 1000 are provided with respect to an organizer node of a render setup graph. Reference is made to FIG. 4 depicting a render setup graph 200 including organizer nodes 210A and 210B. A more detailed depiction of the organizer node 210A is provided in FIG. 5.

[0040] In operation 1002, a dependency graph is accessed. FIG. 4 depicts one type of dependency graph, namely a render setup graph 200. As discussed above, the render setup graph 200 includes a plurality of nodes 204A-E, 210A-B, and 208A-C that are interconnected using node wires 206. The nodes 204A-E may associate one or more visual effects (e.g., lighting or materials) with a set of assets representing computer-generated objects in the scene. The assets and associated visual effects are passed between the interconnected nodes via the node wires 206. As shown in FIG. 4, one or more of the nodes may be an organizer node 210A-B for partitioning the assets into one or more subset asset partitions.

[0041] As described above with respect to FIG. 4, organizer nodes 210A-B are used to select or partition the set of assets in to smaller sets that are used as input to downstream elements of the render setup graphs. Using one or more organizer nodes, a render setup

graph can be configured to define multiple rendering passes, each rendering pass represented by a respective render node 208A-C.

[0042] FIG. 5 depicts a schematic representation of the organizer node 210A. In a typical implementation, a set of assets is identified by the input 211 of the organizer node 210A. The set of assets may be passed to the input 211 by one or more upstream nodes at input 211. The organizer node 210A partitions the input set of assets into multiple subsets of partition assets by evaluating one or more predicate tests 214A-D. Each predicate test 214A-D may include one or more logical operations that can be used to identify or select a subset of partition assets. The subset of partition assets may then be provided on a respective output 212A-D of the organizer node 210A as a partition and connected to a downstream element using a node wire 206.

[0043] One exemplary predicate test may include a logical operation for identifying assets having asset geometry that satisfies a geometry-based logical operation. For example, this exemplary predicate test may identify assets that have curved surface geometry. With reference to the computer-generated objects associated with scene 100 depicted in FIGS. 2 and 3, this predicate test may be used to identify assets that represent the sphere 152, cylinder 153, and cone 153, which each have curved surface geometry.

[0044] In a typical implementation, the assets that do not satisfy a preceding predicate test (e.g., 214A) serve as the input to a subsequent predicate test (e.g., 214B). Typically, each predicate test receives those assets that do not satisfy any preceding predicate tests, and performs a new logical operation on those assets. If an organizer node 210A is configured in this way, the set of assets at the input 211 may be partitioned into five unique and non-overlapping subsets of assets on corresponding outputs 212A-E.

[0045] To ensure that every predicate test of every organizer node can be evaluated, all of the asset information associated with a scene can be loaded and used as input to the render setup graph. Evaluating the render setup graph using a complete set asset information (including asset geometry) may result in a perfect culling of the assets for each rendering node. However, this approach also requires that all of the asset information be loaded into memory, which, as discussed above, consumes a nontrivial amount of computer resources.

[0046] However, using process 1000, the organizer nodes of the render setup graph can be configured to relax the condition that every subset of assets at an output be unique and non-

overlapping. As a result, the organizer nodes can be used to produce a partially culled but potentially over-inclusive (conservative) set of assets can be identified using the render setup graph without having all of the asset information loaded into memory.

[0047] In operation 1004, an input set of assets is accessed at an input of the organizer node. With reference to the organizer node 210A in FIG. 5, an input set of assets at input 211 may be accessed. In this example, the input set of assets is passed to the organizer node 210A from an upstream node via one or more node wires 206. Typically, the input set of assets does not include complete asset information for every asset associated with a scene. For example, the input set of assets may not include geometry information for at least some of the assets associated with the scene. In some cases, the input set of assets includes complete information for at least some, but not all of the assets associated with the scene. The information that is received in operation 1004 typically depends on the asset information that has been previously loaded into workstation memory.

[0048] In many cases, at least some information about all of the assets associated with a scene is available for access by the organizer node 210A in operation 1004. For example, one or more asset attributes for all of the assets associated with a scene may be received by the organizer node 210A. That is, even if all of the assets associated with a scene are not loaded in workstation memory, at least some of the attributes of every asset associated with a scene may be loaded into workstation memory before operation 1004 is performed.

[0049] In general, loading asset attributes consumes significantly less resources than loading the entire asset, which typically includes surface geometry of the computer-generated object. For example, asset attributes, including asset name, location, and bounding box, are typically loaded for every asset associated with a scene. An asset name is typically a unique character string that identifies the asset, a location is typically a three-dimensional coordinate value of the geometric center of the computer-generated object within the scene, and a bounding box is typically a cuboid that defines the volumetric envelope occupied by the computer-generated object in the scene. These asset attributes are typically loaded into local memory because they require a comparatively small amount of computing resources to load and store.

[0050] In operation 1006, it is determined if a predicate test can be evaluated using the input set of assets. As described above with respect to FIG. 5, a predicate test 214A-D

typically includes one or more logical operations that can be used to identify a subset of partition assets by performing a logical operation. Operation 1006 may be performed for each predicate test 214A-D of the organizer node 210A.

[0051] Operation 1006 may be performed by examining the information that is required to perform the logical operation of the predicate test. If the information that is required to perform the logical operation is provided as asset information in operation 1004, the predicate test can typically be evaluated. If the information that is required to perform the logical operation has not been provided as input, or has not been loaded into workstation memory, the predicate test typically cannot be evaluated.

[0052] In a first example, a first exemplary predicate test may be configured to identify assets that have geometry with a curved surface. This first exemplary predicate test can be evaluated if, for example, the geometry of the set of assets is provided as input in operation 1004. However, if the geometry of the set of assets is not provided as input, because, for example, the geometry has not been loaded into the workstation memory, this first exemplary predicate test cannot be evaluated.

[0053] In a second example, a second exemplary predicate test may be configured to identify assets that are within the frustum of the camera. This second exemplary predicate test can be evaluated if the location of each of the assets within the scene is provided as input in operation 1004. The location of the assets can be determined if, for example, the location and bounding box asset attributes for the set of assets received at operation 1004 are known. As discussed above, location and bounding box asset attributes are typically loaded and stored for every asset associated with the scene and are likely to be accessed in operation 1004. If this is the case, and the location and bounding box information is available for access, this second exemplary predicate test can be evaluated.

[0054] In operation 1008, if the predicate test can be evaluated using the input set of assets, a subset of partition assets is identified using the first predicate test. In other words, if there is enough information about the assets to evaluate the predicate test, then identify those assets that satisfy the predicate test. This operation is in accordance with the normal operation of the predicate tests 214A-D of the organizer node 210A, as discussed above with respect to FIG. 5.

[0055] With reference again to the first example provided above, a first exemplary predicate test may be configured to identify assets that have geometry with a curved surface. In this case, the first exemplary predicate test may evaluate the curvature of the asset geometry of each asset of the set of assets. A geometric-based logical operation that tests for geometry having a non-zero curvature can be used to identify the relevant partition assets. With reference to FIG. 1, this first exemplary predicate test may be used to identify those assets that represent the sphere 152, cylinder 153, cone 153, and teapot 151, which all include geometry having a curved surface.

[0056] With reference to the second example provided above, a second exemplary predicate test may be configured to identify assets that are within the frustum of the camera. In this case, the second exemplary predicate test may be evaluated by determining the geometry of the camera frustum identifying assets having a location and bounding box that intersect the frustum geometry. With reference to FIGS. 1-3, this second exemplary predicate test may be used to identify assets that represent the teapot 151, cone 153, and cube 154, which have locations and bounding boxes that intersect the frustum of the virtual camera 104. See, for example, FIG. 3 depicting a rendered image 101 of the scene which shows these objects in view of the camera 104.

[0057] In operation 1010, the subset of partition assets is passed to the next node in the render setup graph. With reference to FIGS. 4 and 5, the assets identified by, for example, partition test 214A is passed to the next downstream node 204C via node wire 206. Visual effects that have already been associated with the set of partition assets by, for example, upstream node 204A, are also passed to the next node in the render setup graph. Because the subset of partition assets are a subset of the input set of assets, the assets have been at least partially culled for downstream node operations.

[0058] As discussed above with respect to FIG. 5 the assets that do not satisfy the predicate test are passed to either the next predicate test in the organizer node or an output of the organizer node. With reference to FIG. 5, assets that do not satisfy predicate test 214A are passed to the next predicate test 214B. In the case that the predicate test is the last test in the series (e.g., predicate test 214D), then the assets that do not satisfy the predicate test are passed to an output (e.g., output 212E) of the organizer node. As shown in FIG. 5, the last output 212E may not be connected to a downstream node by a node wire. In general, any

output of an organizer node may or may not be connected (via a node wire) to a downstream node.

[0059] In operation 1012, if the predicate test cannot be evaluated using the input set of assets, a conservative set of assets is passed to the next node in the render setup graph. To avoid dropping an asset that may be necessary for the rendering pass, the conservative set of assets is the same as the input set of assets. In other words, if there is not enough information about the assets to evaluate the predicate test, then the predicate test is skipped and the set of assets received at the input are reproduced on the respective output of the organizer node as a conservative set of assets and passed to the next node in the render setup graph.

[0060] With reference again to the first example provided above, a first exemplary predicate test may be configured to identify assets that have geometry with a curved surface. In this case, the first exemplary predicate test may not be evaluated if, for example, the asset geometry is not included in the set of assets received in operation 1004. As a result, the entire set of assets that is received at input in operation 1004 is passed to the next node in the render setup graph as a conservative set of assets.

[0061] With reference to FIG. 5, if predicate test 214A cannot be evaluated, the assets that are received at input 211 are provided as a conservative set of assets to an output 212A that corresponds to predicate test 214A. Furthermore, the conservative set of assets (the entire set of assets received by the predicate test 214A) is also passed down to the next predicate test in the series, 214B. In the case of the last predicate test 214D, conservative set of assets (the entire set of assets received by the preceding predicate test 214C) is passed to the output 212E.

[0062] Process 1000 is typically performed for every predicate test in an organizer node. In general, when evaluation of a predicate test is possible, a set of partition assets are identified and passed on through the render setup graph. When evaluation of a predicate test is not possible, the partition test is skipped and the set of assets that are received as an input are passed on through the render setup graph. This process may also be repeated for every organizer node wired to the render node associated with the render pass that is being computed.

[0063] Typically, performing process 1000 for an entire pass of the render setup graph results in a partially-culled, but potentially over-inclusive set of assets at the render node.

That is, that the set of assets is at least partially culled by one or more partitions in the graph that can be evaluated. The set of asset may also be over-inclusive or conservative due to one or more partitions that cannot be evaluated. Thus, using process 1000, the number of assets may be reduced without eliminating assets that may be necessary for the rendering pass.

2. Evaluating a Dependency to Render an Image

[0064] Process 1000 can be used within a broader process to render an image of a scene comprising one or more computer-generated objects. FIG. 7 depicts an exemplary process 1100 for computing a rendered image by performing multiple passes through a dependency graph. More specifically, a first pass of the dependency graph is performed to identify a potentially over-inclusive set of assets necessary for the rendered image. A second pass of the dependency graph is then performed after retrieving some or all of the asset information for assets identified by the first pass. Additional passes can be performed to further refine the set of assets, and then used to retrieve some or all of the asset information for the refined set of assets. Eventually, a rendered image can be calculated based on the last pass of the dependency graph.

[0065] Process 1100, described below, describes a two-pass iteration for refining a set of assets for a rendering pass. However, in other implementations, more than two passes can be performed, with some, but not necessarily all, asset information loaded after each pass through the dependency graph. Performing multiple-passes to refine the set of assets, such as described in process 1100, in conjunction with process 1000, described above, may result in a computational savings by reducing the amount of asset information retrieved from a centralized database.

[0066] In operation 1102, a dependency graph is accessed. As described above with respect to FIG. 4, an exemplary dependency graph is render setup graph 200 including a plurality of nodes 204A-E, 210A-B, and 208A-C that are interconnected using node wires 206. Operation 1102 is similar to operation 1002 discussed above with respect to FIG. 6.

[0067] In operation 1104, a render node is identified. With reference to FIG. 4, one of the render nodes 208A, 208B, or 208C within the render setup graph 200 is identified. The render node may be identified by, for example, selecting a particular render pass of the render setup graph. As discussed above, each render node 208A-C may represent a different render

pass, in which a lighting condition is applied to a specific set of assets (representing computer-generated object in the scene).

[0068] In operation 1006, the nodes that are upstream to the render node are identified. With reference to FIG. 4, multiple nodes 204A, 210A, 204C are upstream of the render node 208A. The upstream nodes define the lighting effects and assets that will be implemented in the render pass. The upstream nodes may be identified by, for example, tracing the node wires from the render node back to the input of the dependency graph.

[0069] In operation 1108, a first pass of the dependency graph is performed through the identified nodes. In this operation, an input set of assets is provided as input to the series of nodes identified in operation 1106. As described above, the input set of assets typically does not include all of the asset information for the set of assets. For example, memory-intensive aspects, such as the asset geometry, may not be loaded and provided as input to the dependency graph to conserve computing resources. As also discussed above with respect to operation 1004 of process 1000, there are typically at least some asset attributes that are loaded for all of the assets associated with the scene. The asset attributes that are loaded on the workstation computer are typically provided as input to the first pass of the dependency graph.

[0070] When performing operation 1108, process 1000 may be implemented for each predicate test that is encountered in the organizer nodes of the render setup graph. As discussed above in the preceding section, implementing process 1000 throughout the predicate test in a dependency graph may result in the identification of a partially culled, but potentially over-inclusive set of assets.

[0071] In operation 1110, some or all of the asset information for the assets that were identified in the first pass are loaded into memory. In one example, the set of assets identified in operation 1008 is used to download all of the asset information associated for those assets from the centralized database storage. As previously mentioned, less than all of the asset information can also be loaded and multiple passes of the dependency graph be performed to refine the set of assets for the render pass.

[0072] Because the assets identified in the first pass (operation 1108) are over-inclusive, the set of assets can be used to identify all of the asset information that may be necessary for performing the selected render pass. Additionally, because the set of assets may also be

partially culled, retrieving the asset information for the identified assets typically results in a computational savings as compared to the download of all of the asset information associated with a scene.

[0073] In operation 1112, a second pass of the dependency graph is performed. In this operation, the dependency graph is evaluated using the asset information obtained in operation 1110. The set of assets obtained in operation 1110 is potentially over-inclusive. Therefore, if all of the asset information is obtained for the assets identified in the first pass, the dependency graph, each of the predicate tests in the render pass can be evaluated and a perfectly culled set of assets can be obtained.

[0074] As previously mentioned, less than all of the asset information may be retrieved for a second pass. If the dependency graph can be fully evaluated on the second pass, then no further passes are necessary. However, if the predicate tests cannot be fully evaluated in the second pass, additional asset information may be retrieved and the dependency graph reevaluated. This process can be repeated as often as necessary to obtain a perfectly culled set of assets.

[0075] In operation 1114, a rendered image is computed. After the second (or subsequent) pass of the dependency graph is performed, a perfectly culled set of assets and the associated visual effects are provided to the respective render node. This information can then be passed to a rendering process to compute a rendered image. Computing a rendered image may include performing one or more processes on the assets using the selected visual effect(s) based on the dependency graph. The external processes may include any process, such as rendering, shading, rasterizing, or the like that the animator, lighter, or other user may want to perform on the objects used by the dependency graph. An exemplary rendered image is depicted in FIG. 3.

3. Workstation Computer System

[0076] FIG. 8 depicts an exemplary workstation computer system 2000 that can be used to implement the render setup graph and techniques discussed above. The render setup graph can be implemented, for example, in either hardware or in software stored on a non-transitory computer-readable storage medium. The system can be configured to generate, modify, and evaluate the render setup graph to configure and manage lighting configuration data as well as external processes used to render a computer-generated image. The system can be further

configured to receive input from a user and to display graphics, an image, or scene of an animation based on the render setup graph.

[0077] The workstation computer system 2000 can be configured to receive user input from an input device 2020. The input device 2020 can be any device that receives input from the user and transmits it to the workstation computer system 2000. For example, the input device 2020 can be a keyboard, a mouse, a tablet, a stylus, or the like. Those skilled in the art will recognize that other types of input devices can also be used.

[0078] The workstation computer system 2000 can be configured to output graphics, images, or animation to a display device 2030. The display device 2030 can include any device that receives data from the workstation computer system and presents the data to the user. For example, the display device 2030 may include a liquid crystal display, a set of light-emitting diodes, a projector, or the like. Those skilled in the art will recognize that other types of output devices can also be used.

[0079] The workstation computer system 2000 may further include a central processing unit 2002. The central processing unit may include one or more processing cores. The central processing unit 2002 may be coupled to and able to communicate with the input device 2020. Although the workstation computer system 2000 is illustrated with one central processing unit 2002, the workstation computer system 2000 may include multiple processing units. The workstation computer system 2000 may also include a graphics processing unit 2004. The graphics processing unit 2004 may be dedicated to processing graphics-related data. The graphics processing unit 2004 may include a single processing core or multiple processing cores. Although the workstation computer system 2000 is illustrated with one graphics processing unit 2004, the workstation computer system 2000 may include a plurality of graphics processing units. The central processing unit 2002 and/or the graphics processing unit 2004 may be coupled to and able to communicate data to the output device 2030.

[0080] In one example, the workstation computer system 2000 may include one or more processors and instructions stored in a non-transitory computer-readable storage medium, such as a memory or storage device, that when executed by the one or more processors, perform animation rendering using a render setup graph, as described above. In the context of the embodiments described herein, a “non-transitory computer-readable storage medium” can be any medium that can contain or store the program for use by or in connection with the

instruction execution system, apparatus, or device. The non-transitory computer readable storage medium can include, but is not limited to, an electronic, magnetic, optical, electromagnetic, infrared, or semiconductor system, apparatus or device, a portable computer diskette (magnetic), a random access memory (RAM) (magnetic), a read-only memory (ROM) (magnetic), an erasable programmable read-only memory (EPROM) (magnetic), a portable optical disc such a CD, CD-R, CD-RW, DVD, DVD-R, or DVD-RW, or flash memory such as compact flash cards, secured digital cards, USB memory devices, memory sticks, and the like.

[0081] The workstation computer system 2000 may include volatile memory 2006, which is a non-transitory computer-readable storage medium, in communication with the central processing unit 2002. The volatile memory 2006 may include, for example, random access memory, such as dynamic random access memory or static random access memory, or any other type of volatile memory. The volatile memory 2006 may be used to store data or instructions during the operation of the workstation computer system 2000. Those skilled in the art will recognize that other types of volatile memory can also be used.

[0082] The workstation computer system 2000 may also include non-volatile memory 2008, which is a non-transitory computer-readable storage medium, in communication with the central processing unit 2002. The non-volatile memory 2008 may include flash memory, hard disks, magnetic storage devices, read-only memory, or the like. The non-volatile memory 2008 may be used to store animation data, render setup graph data, computer instructions, or any other information. Those skilled in the art will recognize that other types of non-volatile memory can also be used.

[0083] The workstation computer system 2000 is not limited to the devices, configurations, and functionalities described above. For example, although a single volatile memory 2006, non-volatile memory 2008, central processing unit 2002, graphics processing unit 2004, input device 2020, and output device 2030 are illustrated, a plurality of any of these devices can be implemented internal or external to the workstation computer system 2000. In addition, the workstation computer system 2000 may include a network access device for accessing information on a network, such as an internal network or the Internet. Those skilled in the art will recognize that other configurations of the workstation computer system 2000 can be used.

[0084] Various exemplary embodiments are described herein. Reference is made to these examples in a non-limiting sense. They are provided to illustrate more broadly applicable aspects of the disclosed technology. Various changes may be made and equivalents may be substituted without departing from the true spirit and scope of the various embodiments. In addition, many modifications may be made to adapt a particular situation, material, composition of matter, process, process act(s), or step(s) to the objective(s), spirit, or scope of the various embodiments. Further, as will be appreciated by those with skill in the art, each of the individual variations described and illustrated herein has discrete components and features that may be readily separated from or combined with the features of any of the other several embodiments without departing from the scope or spirit of the various embodiments.

CLAIMS

What is claimed is:

1. A computer-implemented method for partitioning a set of assets, wherein each asset represents a computer-generated object associated with a computer-generated scene, the method comprising:
 - accessing a dependency graph, the dependency graph comprising a plurality of interconnected nodes, wherein one of the interconnected nodes includes an organizer node having at least one predicate test;
 - accessing the set of assets identified by an input of the predicate test of the organizer node;
 - determining if the at least one predicate test can be evaluated using the set of assets;
 - if the at least one predicate test can be evaluated,
 - performing a logical operation on the set of assets to identify one or more partition assets, wherein the one or more partition assets are a subset of the set of assets, and
 - passing the one or more partition assets at an output of the organizer node to a next node in the dependency graph; and
 - if the at least one predicate test cannot be evaluated,
 - passing a conservative set of assets at the output of the organizer node to the next node, wherein the conservative set of assets is the same set of assets identified by the input of the predicate test.
2. The computer-implemented method of claim 1, wherein the set of assets includes asset geometry and asset attributes for one or more assets of the set of assets.
3. The computer-implemented method of claim 1, wherein the set of assets does not include asset geometry for one or more assets of the set of assets.
4. The computer-implemented method of claim 1, wherein the at least one predicate test can be evaluated if a first asset attribute is required for the logical operation and the first asset attribute is able to be accessed at the input to the predicate test.

5. The computer-implemented method of claim 1, wherein the at least one predicate test cannot be evaluated if a first asset attribute is required for the logical operation and the first asset attribute is not able to be accessed at the input to the predicate test.

6. The computer-implemented method of claim 1, wherein the at least one predicate test cannot be evaluated if the asset geometry is required for the logical operation and the asset geometry is able to be accessed at the input to the predicate test.

7. The computer-implemented method of claim 1, wherein the at least one predicate test cannot be evaluated because at least some information for one or more assets of the set of assets is not loaded in computer memory.

8. The computer-implemented method of claim 1, wherein one or more nodes of the plurality of nodes assigns a visual effect to the set of assets.

9. The computer-implemented method of claim 1, wherein the visual effect is a virtual light source.

10. The computer-implemented method of claim 1, further comprising:
passing the one or more partition assets to a render node of the dependency graph; and

computing a rendered image of the computer-generated objects represented by the one or more partition assets using the render node.

11. A computer-implemented method for partitioning a set of assets, wherein each asset represents a computer-generated object associated with a computer-generated scene, the method comprising:

accessing a dependency graph, the dependency graph comprising a plurality of interconnected nodes, including a render node associated with a rendering pass;

identifying a render node of the dependency graph, the render node associated with a rendering pass of the computer-generated scene;

identifying a set of nodes of the dependency graph that are connected to the render node and are upstream from the render node;

performing a first pass through the dependency graph through the set of nodes to identify a set of partition assets;
loading asset information for the set of partition assets;
performing a second pass through the dependency graph through the set of nodes using the asset information for the set of partition assets; and
computing a rendered image based on the second pass through the dependency graph.

12. The computer-implemented method of claim 11, wherein performing a first pass through the dependency graph comprises:

accessing the dependency graph, wherein one of the interconnected nodes includes an organizer node having at least one predicate test;

accessing the set of assets identified by an input of the predicate test of the organizer node;

determining if the at least one predicate test can be evaluated using the set of assets;

if the at least one predicate test can be evaluated,

performing a logical operation on the set of assets to identify one or more partition assets, wherein the one or more partition assets are a subset of the set of assets, and

passing the one or more partition assets at an output of the organizer node to a next node in the dependency graph; and

if the at least one predicate test cannot be evaluated,

passing a conservative set of assets at the output of the organizer node to the next node, wherein the conservative set of assets is the same set of assets identified by the input of the predicate test.

13. A non-transitory computer-readable storage medium comprising computer-executable instructions for partitioning a set of assets, wherein each asset represents a computer-generated object associated with a computer-generated scene, the computer-executable instructions comprising instructions for:

accessing a dependency graph, the dependency graph comprising a plurality of interconnected nodes, wherein one of the interconnected nodes includes an organizer node having at least one predicate test;

accessing the set of assets identified by an input of the predicate test of the organizer node;

determining if the at least one predicate test can be evaluated using the set of assets;

if the at least one predicate test can be evaluated,

performing a logical operation on the set of assets to identify one or more partition assets, wherein the one or more partition assets are a subset of the set of assets, and

passing the one or more partition assets at an output of the organizer node to a next node in the dependency graph; and

if the at least one predicate test cannot be evaluated,

passing a conservative set of assets at the output of the organizer node to the next node, wherein the conservative set of assets is the same set of assets identified by the input of the predicate test.

14. The computer-readable storage medium of claim 13, wherein the set of assets includes asset geometry and asset attributes for one or more assets of the set of assets.

15. The computer-readable storage medium of claim 13, wherein the set of assets does not include asset geometry for one or more assets of the set of assets.

16. The computer-readable storage medium of claim 13, further comprising instructions for:

passing the one or more partition assets to a render node of the dependency graph; and

computing a rendered image of the computer-generated objects represented by the one or more partition assets using the render node.

17. A non-transitory computer-readable storage medium comprising computer-executable instructions for partitioning a set of assets, wherein each asset represents a computer-generated object associated with a computer-generated scene, the computer-executable instructions comprising instructions for:

accessing a dependency graph, the dependency graph comprising a plurality of interconnected nodes, including a render node associated with a rendering pass;

identifying a render node of the dependency graph, the render node associated with a rendering pass of the computer-generated scene;

identifying a set of nodes of the dependency graph that are connected to the render node and are upstream from the render node;

performing a first pass through the dependency graph through the set of nodes to identify a set of partition assets;

loading asset information for the set of partition assets;

performing a second pass through the dependency graph through the set of nodes using the asset information for the set of partition assets; and

computing a rendered image based on the second pass through the dependency graph.

18. An apparatus for performing graphics processing, the apparatus comprising:

a memory configured to store data; and

a computer processor configured to:

access a dependency graph, the dependency graph comprising a plurality of interconnected nodes, wherein one of the interconnected nodes includes an organizer node having at least one predicate test;

access the set of assets identified by an input of the predicate test of the organizer node;

determine if the at least one predicate test can be evaluated using the set of assets;

if the at least one predicate test can be evaluated,

perform a logical operation on the set of assets to identify one or more partition assets, wherein the one or more partition assets are a subset of the set of assets, and

pass the one or more partition assets at an output of the organizer node to a next node in the dependency graph; and

if the at least one predicate test cannot be evaluated,

pass a conservative set of assets at the output of the organizer node to the next node, wherein the conservative set of assets is the same set of assets identified by the input of the predicate test.

19. The apparatus of claim 18, wherein the computer processor is further configured to:
- pass the one or more partition assets to a render node of the dependency graph;
 - and
 - compute a rendered image of the computer-generated objects represented by the one or more partition assets using the render node.

20. An apparatus for performing graphics processing, the apparatus comprising:

- a memory configured to store data; and

- a computer processor configured to:

 - access a dependency graph, the dependency graph comprising a plurality of interconnected nodes, including a render node associated with a rendering pass;

 - identify a render node of the dependency graph, the render node associated with a rendering pass of the computer-generated scene;

 - identify a set of nodes of the dependency graph that are connected to the render node and are upstream from the render node;

 - perform a first pass through the dependency graph through the set of nodes to identify a set of partition assets;

 - load asset information for the set of partition assets;

 - perform a second pass through the dependency graph through the set of nodes using the asset information for the set of partition assets; and

 - compute a rendered image based on the second pass through the dependency graph.




FIG. 1




FIG. 2




FIG. 3




FIG. 4




FIG. 5





FIG. 6

**FIG. 7**




FIG. 8