

CONTEXT MODELING OF OCCUPANCY CODING FOR POINT CLOUD CODING

CROSS REFERENCE TO RELATED APPLICATIONS

[0001] This present application claims the benefit of priority to U.S. Patent Application No. 17/231,695, “CONTEXT MODELING OF OCCUPANCY CODING FOR POINT CLOUD CODING” filed on April 15, 2021, which claims priority based on U.S. Provisional Application Nos. 63/034,113, “UPDATES ON CONTEXT MODELING OF OCCUPANCY CODING FOR POINT CLOUD CODING” filed on June 3, 2020 and 63/066,099, “UPDATES ON PARENT NODE LEVEL CONTEXT MODELING OF OCCUPANCY CODING FOR POINT CLOUD CODING” filed on August 14, 2020, the entirety of which are incorporated by reference herein.

BACKGROUND

[0002] This disclosure relates generally to field of data processing, and more particularly to point clouds.

[0003] Point Cloud has been widely used in recent years. For example, it is used in autonomous driving vehicles for object detection and localization; it is also used in geographic information systems (GIS) for mapping, and used in cultural heritage to visualize and archive cultural heritage objects and collections, etc. Point clouds contain a set of high dimensional points, typically of three dimensional (3D), each including 3D position information and additional attributes such as color, reflectance, etc. They can be captured using multiple cameras and depth sensors, or Lidar in various setups, and may be made up of thousands up to billions of points to realistically represent the original scenes. Compression technologies are needed to reduce the amount of data required to represent a point cloud for faster transmission or reduction of storage. ISO/IEC MPEG (JTC 1/SC 29/WG 11) has created an ad-hoc group (MPEG-PCC) to standardize the compression techniques for static or dynamic point clouds.

SUMMARY

[0004] Embodiments relate to a method, system, and computer readable medium for decoding point cloud data. According to one aspect, a method for decoding point cloud data is provided. The method may include receiving data corresponding to a point cloud. A number of contexts associated with the received data is reduced based on occupancy data corresponding to one or more parent nodes and one or more child nodes within the received data. The data corresponding to the

point cloud is decoded based on the reduced number of contexts.

[0005] According to another aspect, a computer system for decoding point cloud data is provided. The computer system may include one or more processors, one or more computer-readable memories, one or more computer-readable tangible storage devices, and program instructions stored on at least one of the one or more storage devices for execution by at least one of the one or more processors via at least one of the one or more memories, whereby the computer system is capable of performing a method. The method may include receiving data corresponding to a point cloud. A number of contexts associated with the received data is reduced based on occupancy data corresponding to one or more parent nodes and one or more child nodes within the received data. The data corresponding to the point cloud is decoded based on the reduced number of contexts.

[0006] According to yet another aspect, a computer readable medium for decoding point cloud data is provided. The computer readable medium may include one or more computer-readable storage devices and program instructions stored on at least one of the one or more tangible storage devices, the program instructions executable by a processor. The program instructions are executable by a processor for performing a method that may accordingly include receiving data corresponding to a point cloud. A number of contexts associated with the received data is reduced based on occupancy data corresponding to one or more parent nodes and one or more child nodes within the received data. The data corresponding to the point cloud is decoded based on the reduced number of contexts.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] These and other objects, features and advantages will become apparent from the following detailed description of illustrative embodiments, which is to be read in connection with the accompanying drawings. The various features of the drawings are not to scale as the illustrations are for clarity in facilitating the understanding of one skilled in the art in conjunction with the detailed description. In the drawings:

FIG. 1 illustrates a networked computer environment according to at least one embodiment;

FIG. 2A is a diagram of an octree structure for point cloud data, according to at least one embodiment;

FIG. 2B is a diagram of an octree partition for point cloud data, according to at least one embodiment;

FIG. 2C is a block diagram of parent-level node neighbors for point cloud data, according to at least one embodiment;

FIG. 2D is a block diagram of parent-node-level contexts for point cloud data, according to at least one embodiment;

FIG. 2E is a block diagram of parent-level neighbors of a current coded node for point cloud data, according to at least one embodiment;

FIG. 2F is a block diagram of parent-node-level contexts for point cloud data, according to at least one embodiment;

FIG. 2G is a block diagram of parent-level neighbors of a current coded node for point cloud data, according to at least one embodiment;

FIG. 2H is a block diagram of child-node-level neighbors for point cloud data, according to at least one embodiment

FIG. 3A is a syntax element for signaling a reduction in point cloud data contexts, according to at least one embodiment;

FIG. 3B is a hash table for coded point cloud geometry information, according to at least one embodiment;

FIG. 4 is an operational flowchart illustrating the steps carried out by a program that decodes point cloud data, according to at least one embodiment;

FIG. 5 is a block diagram of internal and external components of computers and servers depicted in FIG. 1 according to at least one embodiment;

FIG. 6 is a block diagram of an illustrative cloud computing environment including the computer system depicted in FIG. 1, according to at least one embodiment; and

FIG. 7 is a block diagram of functional layers of the illustrative cloud computing environment of FIG. 6, according to at least one embodiment.

DETAILED DESCRIPTION

[0008] Detailed embodiments of the claimed structures and methods are disclosed herein; however, it can be understood that the disclosed embodiments are merely illustrative of the claimed structures and methods that may be embodied in various forms. Those structures and methods may, however, be embodied in many different forms and should not be construed as limited to the exemplary embodiments set forth herein. Rather, these exemplary embodiments are provided so that this disclosure will be thorough and complete and will fully convey the scope to those skilled in the art. In the description, details of well-known features and techniques may be omitted to avoid unnecessarily obscuring the presented embodiments.

[0009] Embodiments relate generally to the field of data processing, and more particularly to point clouds. The following described exemplary embodiments provide a system, method and computer program to, among other things, reduce a number of contexts associated with point cloud data. Therefore, some embodiments have the capacity to improve the field of computing by allowing for improved point cloud compression and decompression by considering a relevant subset of neighboring nodes from among all neighboring nodes of a current node in point cloud data.

[0010] As previously described, point cloud has been widely used in recent years. For example, it is used in autonomous driving vehicles for object detection and localization; it is also used in geographic information systems (GIS) for mapping, and used in cultural heritage to visualize and archive cultural heritage objects and collections, etc. Point clouds contain a set of high dimensional points, typically of three dimensional (3D), each including 3D position information and additional attributes such as color, reflectance, etc. They can be captured using multiple cameras and depth sensors, or Lidar in various setups, and may be made up of thousands up to billions of points to realistically represent the original scenes. Compression technologies are needed to reduce the amount of data required to represent a point cloud for faster transmission or

reduction of storage. ISO/IEC MPEG (JTC 1/SC 29/WG 11) has created an ad-hoc group (MPEG-PCC) to standardize the compression techniques for static or dynamic point clouds.

[0011] However, the parent-node-level context modeling of occupancy coding restricts itself to accessing up to 6 neighboring coded nodes. This may lead to a relatively small number of contexts. Accordingly, only looking at 6 nearest neighbors will largely restrain the range of receptive field and thus limit the performance gain from it. Moreover, the context modeling of occupancy coding may result in a large number of contexts. However, in terms of hardware implementation, a very large number of contexts is not preferable since they consume a large amount of memory. It may be advantageous, therefore, to improve the performance of parent-node-level context while keeping the number of contexts small by reducing the number of contexts presented.

[0012] Aspects are described herein with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems), and computer readable media according to the various embodiments. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer readable program instructions.

[0013] The following described exemplary embodiments provide a system, method and computer program that compresses and decompresses point cloud data based on reducing an expanded set of contexts associated with the point cloud data. Referring now to FIG. 1, a functional block diagram of a networked computer environment illustrating a point cloud compression system 100 (hereinafter “system”) for compressing and decompressing point cloud data. It should be appreciated that FIG. 1 provides only an illustration of one implementation and does not imply any limitations with regard to the environments in which different embodiments may be implemented. Many modifications to the depicted environments may be made based on design and implementation requirements.

[0014] The system 100 may include a computer 102 and a server computer 114. The computer 102 may communicate with the server computer 114 via a communication network 110 (hereinafter “network”). The computer 102 may include a processor 104 and a software program 108 that is stored on a data storage device 106 and is enabled to interface with a user and communicate with the server computer 114. As will be discussed below with reference to FIG. 5

the computer 102 may include internal components 800A and external components 900A, respectively, and the server computer 114 may include internal components 800B and external components 900B, respectively. The computer 102 may be, for example, a mobile device, a telephone, a personal digital assistant, a netbook, a laptop computer, a tablet computer, a desktop computer, or any type of computing devices capable of running a program, accessing a network, and accessing a database.

[0015] The server computer 114 may also operate in a cloud computing service model, such as Software as a Service (SaaS), Platform as a Service (PaaS), or Infrastructure as a Service (IaaS), as discussed below with respect to FIGS. 6 and 7. The server computer 114 may also be located in a cloud computing deployment model, such as a private cloud, community cloud, public cloud, or hybrid cloud.

[0016] The server computer 114, which may be used for point cloud compression and decompression is enabled to run a Point Cloud Compression Program 116 (hereinafter “program”) that may interact with a database 112. The Point Cloud Compression Program method is explained in more detail below with respect to FIG. 4. In one embodiment, the computer 102 may operate as an input device including a user interface while the program 116 may run primarily on server computer 114. In an alternative embodiment, the program 116 may run primarily on one or more computers 102 while the server computer 114 may be used for processing and storage of data used by the program 116. It should be noted that the program 116 may be a standalone program or may be integrated into a larger point cloud compression and decompression program.

[0017] It should be noted, however, that processing for the program 116 may, in some instances be shared amongst the computers 102 and the server computers 114 in any ratio. In another embodiment, the program 116 may operate on more than one computer, server computer, or some combination of computers and server computers, for example, a plurality of computers 102 communicating across the network 110 with a single server computer 114. In another embodiment, for example, the program 116 may operate on a plurality of server computers 114 communicating across the network 110 with a plurality of client computers. Alternatively, the program may operate on a network server communicating across the network with a server and a plurality of client computers.

[0018] The network 110 may include wired connections, wireless connections, fiber optic connections, or some combination thereof. In general, the network 110 can be any combination of connections and protocols that will support communications between the computer 102 and the server computer 114. The network 110 may include various types of networks, such as, for example, a local area network (LAN), a wide area network (WAN) such as the Internet, a telecommunication network such as the Public Switched Telephone Network (PSTN), a wireless network, a public switched network, a satellite network, a cellular network (e.g., a fifth generation (5G) network, a long-term evolution (LTE) network, a third generation (3G) network, a code division multiple access (CDMA) network, etc.), a public land mobile network (PLMN), a metropolitan area network (MAN), a private network, an ad hoc network, an intranet, a fiber optic-based network, or the like, and/or a combination of these or other types of networks.

[0019] The number and arrangement of devices and networks shown in FIG. 1 are provided as an example. In practice, there may be additional devices and/or networks, fewer devices and/or networks, different devices and/or networks, or differently arranged devices and/or networks than those shown in FIG. 1. Furthermore, two or more devices shown in FIG. 1 may be implemented within a single device, or a single device shown in FIG. 1 may be implemented as multiple, distributed devices. Additionally, or alternatively, a set of devices (e.g., one or more devices) of system 100 may perform one or more functions described as being performed by another set of devices of system 100.

[0020] Referring now to FIG. 2A, a diagram of an octree structure 200A is depicted. In TMC13, if the octree geometry codec is used, the geometry encoding proceeds as follows. First, a cubical axis-aligned bounding box B is defined by two points $(0,0,0)$ and $(2^{M-1}, 2^{M-1}, 2^{M-1})$, where 2^{M-1} defines the size of B and M is specified in the bitstream. The octree structure 200A is then built by recursively subdividing B. At each stage, a cube is subdivided into 8 sub-cubes. An 8-bit code, namely the occupancy code, is then generated by associating a 1-bit value with each sub-cube in order to indicate whether it contains points (i.e., full and has value 1) or not (i.e., empty and has value 0). Only full sub-cubes with a size greater than 1 (i.e., non-voxels) are further subdivided.

[0021] Referring now to FIG. 2B, a diagram of an octree partition 200B is depicted. The octree

partition 200B may include a two-level octree partition 202 and a corresponding occupancy code 204, where cubes and nodes in dark indicate they are occupied by points. The occupancy code 204 of each node is then compressed by an arithmetic encoder. The occupancy code 204 can be denoted as S which is an 8-bit integer, and each bit in S indicates the occupancy status of each child node. Two encoding methods for occupancy code 204 exist in TMC13, i.e., the bit-wise encoding and the byte-wise encoding methods, and the bit-wise encoding is enabled by default. Either way performs arithmetic coding with context modeling to encode the occupancy code 204, where the context status is initialized at the beginning of the whole coding process and is updated during the coding process.

[0022] For bit-wise encoding, eight bins in S are encoded in a certain order where each bin is encoded by referring to the occupancy status of neighboring nodes and child nodes of neighboring nodes, where the neighboring nodes are in the same level of current node. For byte-wise encoding, S is encoded by referring to an adaptive look up table (A-LUT), which keeps track of the N (e.g., 32) most frequent occupancy codes and a cache which keeps track of the last different observed M (e.g., 16) occupancy codes.

[0023] A binary flag indicating whether S is the A-LUT or not is encoded. If S is in the A-LUT, the index in the A-LUT is encoded by using a binary arithmetic encoder. If S is not in the A-LUT, then a binary flag indicating whether S is in the cache or not is encoded. If S is in the cache, then the binary representation of its index is encoded by using a binary arithmetic encoder. Otherwise, if S is not in the cache, then the binary representation of S is encoded by using a binary arithmetic encoder. The decoding process starts by parsing the dimensions of the bounding box B from bitstream. The same octree structure is then built by subdividing B according to the decoded occupancy codes.

[0024] An occupancy code of current node typically has 8 bits, where each bit represents whether its i^{th} child node is occupied or not. When coding the occupancy code of the current node, all the information from neighboring coded nodes can be used for context modeling. The context information can be further grouped in terms of the partition level and distance to current node. Without loss of generality, the context index of the i^{th} child node in current node can be obtained as follows,

$$idx = LUT[i][ctxIdxParent][ctxIdxChild],$$

where LUT is a look-up table of context indices. *ctxIdxParent* and *ctxIdxChild* denote the LUT indices representing the parent-node-level and child-node-level neighboring information.

[0025] Referring now to FIG. 2C, a block diagram 200C of parent-level node neighbors is depicted. According to one or more embodiments, suppose the octree is traversed in the breadth-first order. Six neighboring nodes of current node can be utilized as the parent-node-level context installation. The block in center may be the current node and the other six blocks are may be nearest neighbors where their distance to the current node is 1 block unit. Let P_i denote the occupancy status of the i^{th} neighbor, i.e., $P_i = \begin{cases} 1 & \text{if neighbor } i \text{ is occupied} \\ 0 & \text{if neighbor } i \text{ is not occupied} \end{cases}$, where $i = 0, 1, \dots, 5$. Then the parent-node-level context index can be obtained as follows,

$$ctxIdxParent = \sum_{i=0}^5 2^i P_i.$$

Therefore, if all the neighbors are occupied, the maximum of *ctxIdxParent* is 63, indicating the number of contexts in parent-node-level is 64 in this case.

[0026] Referring now to FIG. 2D, a block diagram 200D of parent-node-level contexts is depicted. According to one or more embodiments, the octree is traversed in the breadth-first order. All the nodes in parent-level occupancy may be coded and they can be used as contexts in current level. The parent-node-level context accessing method may differ for each child node with index i . For example, 8 sub-figures may represent the parent-node-level context accessing method for i^{th} child node ($i = 0, 1, \dots, 7$) of current coded node. In each sub-figure, the current coded node is highlighted in dark gray and the smaller sub-block within the current node is the i^{th} child node ($i = 0, 1, \dots, 7$); each child node has different position relative to the current node. The parent-node-level context accessing method is different for each child node depending on its relative position to the current node.

[0027] Referring now to FIG. 2E, a block diagram 200E of parent-level neighbors of a current coded node is depicted. In one or more embodiments, when coding the i^{th} child node ($i = 0, 1, \dots, 7$), the 7 adjacent nodes to the i^{th} child node in parent-level are utilized as contexts. The 7 adjacent nodes to the i^{th} child node include 3 nodes that share a same face with the child node and 3 nodes that share a same edge and 1 node that shares a same vertex. Since different child nodes

refer to different sets of parent nodes, when coding current node, 26 parent-level neighboring nodes are used in total. Compared to using only 6 parent-level neighboring nodes, utilizing 26 parent-level neighboring nodes leverages much more information.

[0028] Accordingly, the context index of the i^{th} child node in current node can be obtained as:
 $idx = LUT[i][ctxIdxParent][ctxIdxChild],$

$$ctxIdxParent = \sum_{j=0}^6 2^j P_{i,j}$$

where $P_{i,j}$ represents the occupancy of the j^{th} adjacent node in parent-level of the i^{th} child node. $P_{i,j}$ equals to 1 if the j^{th} adjacent node in parent-level of the i^{th} child node is occupied and equals to 0 otherwise, where $i = 0,1, \dots, 7$ and $j = 0,1, \dots, 6$.

[0029] To reduce the context number in parent-node-level, in another embodiment, the parent-node-level context index may be obtained by,

$$ctxIdxParent = \sum_{i=0}^5 P_i,$$

in which case, the maximum of $ctxIdxParent$ is reduced to 6, indicating the number of contexts in parent-node-level is reduced to 7 from 64.

[0030] To further reduce the context number in parent-node-level, in another embodiment, the parent-node-level context is disabled by setting

$$ctxIdxParent = 0.$$

In this case, the parent-node-level neighboring information is not used at all.

[0031] Referring now to FIG. 2F, a block diagram 200F of parent-node-level contexts is depicted. According to one or more embodiments, the number of parent-level neighbors can be changed. For example, each child node may only utilize 6 nearest parent-level neighbors. 8 sub-figures represent the parent-node-level context accessing method for i^{th} child node ($i = 0,1, \dots, 7$) of the current coded node. In each sub-figure, the current coded node is highlighted in dark gray and the smaller sub-block within the current node is the i^{th} child node ($i = 0,1, \dots, 7$); each child node has different position relative to current node. The parent-node-level context accessing method is different for each child node depending on its relative position to the current node. When coding the i^{th} child node ($i = 0,1, \dots, 7$), only 6 adjacent nodes to the i^{th} child node in

parent-level are utilized as contexts. The 6 adjacent nodes to the i^{th} child node may include 3 nodes that share a same face with the child node and 3 nodes that share a same edge. Compared to FIG. 2D, the node that shares a same vertex with child node may not be used in this case.

[0032] Referring to FIG. 2G, a block diagram depicting 200G depicting parent-level neighbors of a current coded node, based on the parent-node-level contexts of FIG. 2F, is depicted. In one or more embodiments, the node that shares a same vertex with child node may not be used. Accordingly, up to 18 parent-level neighboring nodes may be used in total, reduced from the 26 in FIGS. 2D and 2E. Therefore, the context index of the i^{th} child node in current node can be obtained as:

$$idx = LUT[i][ctxIdxParent][ctxIdxChild],$$

$$ctxIdxParent = \sum_{j=0}^5 2^j P_{i,j}$$

where $P_{i,j}$ represents the occupancy of the j^{th} adjacent node in parent-level of the i^{th} child node. $P_{i,j}$ equals to 1 if the j^{th} adjacent node in parent-level of the i^{th} child node is occupied and equals to 0 otherwise, where $i = 0,1, \dots, 7$ and $j = 0,1, \dots, 5$. It may be appreciated that the number of parent-level neighbors may be further reduced by using only a subset of neighbors.

[0033] To further reduce the number of contexts, one can simplify the derivation of $ctxIdxParent$. In one embodiment, each child node refers up to 6 parent-level neighbors as shown in Figure 6, and the parent-level context index can be obtained as:

$$ctxIdxParent = \left(\sum_{j=0}^2 2^j P_{i,j} \right) * 4 + \sum_{j=3}^5 P_{i,j}$$

where $P_{i,j}$ ($j = 0,1,2$) represents the occupancy status of three parent-level nodes that share a same face with the i^{th} child node, and $P_{i,j}$ ($j = 3,4,5$) represents the occupancy status of three parent-level nodes that share a same edge with the i^{th} child node. In this case, the neighbors sharing a face are of higher importance and the context index is calculated by taking their all possible combinations, while the neighbors sharing an edge are of lower importance and the context index is calculated by just counting how many of them are occupied. By this means, the maximum value of $ctxIdxParent$ is reduced from $2^6 - 1$ to $2^5 - 1$, therefore the total number of contexts is reduced by half. Note that other similar simplification methods can be applied to derive parent-level context.

[0034] Referring now to FIG. 2H, a block diagram 200H of child-node-level neighbors is depicted. According to one or more embodiments, the child-node-level neighbors can be defined in many ways. For example, the child-node-level neighbor nodes that are already coded may be classified in terms of their distance to the center node. Without loss of generality, assuming the size of each edge of the child nodes is 1, the ℓ_2 distance to the center is 1, $\sqrt{2}$ and $\sqrt{3}$ from left to right, respectively.

[0035] Let C_i^1 , $C_i^{\sqrt{2}}$ and $C_i^{\sqrt{3}}$ denote the occupancy status of the i^{th} neighbor in child-node-level with ℓ_2 distance equal to 1, $\sqrt{2}$ and $\sqrt{3}$, respectively, i.e., $C_i^* = \begin{cases} 1 & \text{if neighbor } i \text{ is occupied} \\ 0 & \text{if neighbor } i \text{ is not occupied} \end{cases}$. Note that the maximum number of coded child-level neighboring nodes with distance 1, $\sqrt{2}$ and $\sqrt{3}$ may be 3, 9 and 7, respectively.

[0036] Accordingly, the child-node-level context index can be obtained as:

$$ctxIdxChild = \left(\sum_{i=0}^2 2^i C_i^1 \right) * \left(\sum_{i=0}^8 2^i C_i^{\sqrt{2}} \right) * \left(\sum_{i=0}^6 2^i C_i^{\sqrt{3}} \right)$$

Therefore, the maximum number of contexts in child-node-level is 2^{19} , which is too much for applications. To reduce the number of contexts, in one or more embodiments, take partial child-node-level neighbors may be taken into account. This may reduce the maximum context number in child-node-level to $2^7 = 128$.

[0037] In another embodiment, the maximum number of contexts may be further simplified as:

$$ctxIdxChild = \left(\sum_{i=0}^2 2^i C_i^1 \right) * \left(\sum_{i=0}^8 C_i^{\sqrt{2}} + \sum_{i=0}^6 C_i^{\sqrt{3}} \right)$$

where the maximum context number in child-node-level is reduced to $2^3 \times 16 = 128$.

[0038] To further reduce the context number, one may only use the three nearest nodes with ℓ_2 distance 1, and the context index is calculated as:

$$ctxIdxChild = \left(\sum_{i=0}^2 2^i C_i^1 \right),$$

where the maximum context number in child-node-level is reduced to $2^3 = 8$.

[0039] To further reduce the number of contexts, the relationship between the parent-level

context $ctxIdxParent$ and the child-node index i can be decoupled. One can define certain rules to achieve this. In one embodiment, each child node refers up to 6 parent-level neighbors. The context index of the i^{th} child node in current node can be obtained as:

$$idx = LUT[ctxIdxParent][ctxIdxChild],$$

$$ctxIdxParent = \left(\sum_{j=0}^2 2^j P_{i,j} \right) * 4 + \sum_{j=3}^5 P_{i,j}$$

where $P_{i,j}$ ($j = 0,1,2$) represents the occupancy status of three parent-level nodes that share a same face with the i^{th} child node, and $P_{i,j}$ ($j = 3,4,5$) represents the occupancy status of three parent-level nodes that share a same edge with the i^{th} child node. Note that the final context index is independent to child index i . To achieve this, a mapping may be defined for each child node about the ordering of the neighboring parent-level nodes, i.e., $P_{i,j}$. The mapping will be discussed with regard to FIG. 3B. The decoupling reflects the symmetric geometry between each child node and its 6 neighboring parent-level nodes. It may be appreciated that other decoupling methods can be used via defining similar mapping tables in other embodiments.

[0040] According to one or more embodiments, there may be no differentiation between different child nodes when calculating the context in order to reduce the number of contexts. Specifically, the context index of the i^{th} child node in current node can be obtained as follows,

$$idx = LUT[ctxIdxParent][ctxIdxChild],$$

where i is no longer a dimension in the LUT. In other words, all the child nodes are viewed as the same.

[0041] Referring now to FIG. 3A, an syntax element 300A is depicted. In order to achieve adaptivity to different contents and scenarios, different context reduction strategies may be enabled for different cases. Accordingly, a high-level syntax may be signaled in either sequence parameter set or geometry parameter set or slice data header or elsewhere to indicate the context reduction level. In terms of the value of the specified context reduction level, the encoder and decoder can then decide how to utilize context information from neighboring coded nodes. A context reduction level may correspond to a combination scheme of context reduction methods.

[0042] In one or more embodiments, the syntax element 300A may be signaled in geometry parameter set, indicating the context reduction level as shown in the table below. If the context

reduction level is 0, meaning there is no context reduction, the regular context installation is applied. If the context reduction level is 1, a low-level reduction is applied. If the context reduction level is 2, a medium-level reduction is applied. If the context reduction level is 3, a high-level reduction is applied. With the higher number of context reduction level, a smaller number of contexts is used. In this case, different context configurations may be chosen in terms of contents and scenarios. Thus, **gps_context_reduction_level** may specify the context reduction level in geometry parameter set.

[0043] Referring now to FIG. 3B, a hash table 300B is depicted. In one or more embodiments, the coded geometry occupancy information must be saved in cache memory because the later coded nodes need to use this information as context. The hash table 300B may be used to cache occupancy information. For each octree partition depth d , a hash table H_d is maintained, where the key is the Morton code of an octree node in depth d , i.e., $M_i^d = Morton(x_i^d, y_i^d, z_i^d)$, where (x_i^d, y_i^d, z_i^d) is the 3D coordinates of the octree node. Using the Morton code M_i^d as the key, one can access its occupancy value in the hash table H_d . The occupancy value of the octree node i in depth d can be obtained as:

$$OC_i^d = \begin{cases} H_d(M_i^d) & \text{if } H_d(M_i^d) \neq NULL \\ 0 & \text{otherwise} \end{cases}.$$

When encoding/decoding the occupancy value of current node, the occupancy information of neighboring nodes is obtained from the hash table H_d . After encoding/decoding an occupancy value of the current node, the coded occupancy value is then stored in H_d .

[0044] Referring now to FIG. 4, an operational flowchart illustrating the steps of a method 400 carried out by a program that compresses and decompresses point cloud data is depicted.

[0045] At 402, the method 400 may include receiving data corresponding to a point cloud.

[0046] At 404, the method 400 may include reducing a number of contexts associated with the received data based on occupancy data corresponding to one or more parent nodes and one or more child nodes within the received data.

[0047] At 406, the method 400 may include decoding the data corresponding to the point cloud based on the reduced number of contexts.

[0048] It may be appreciated that FIG. 4 provides only an illustration of one implementation and does not imply any limitations with regard to how different embodiments may be implemented. Many modifications to the depicted environments may be made based on design and implementation requirements.

[0049] FIG. 5 is a block diagram 500 of internal and external components of computers depicted in FIG. 1 in accordance with an illustrative embodiment. It should be appreciated that FIG. 5 provides only an illustration of one implementation and does not imply any limitations with regard to the environments in which different embodiments may be implemented. Many modifications to the depicted environments may be made based on design and implementation requirements.

[0050] Computer 102 (FIG. 1) and server computer 114 (FIG. 1) may include respective sets of internal components 800A,B and external components 900A,B illustrated in FIG 5. Each of the sets of internal components 800 include one or more processors 820, one or more computer-readable RAMs 822 and one or more computer-readable ROMs 824 on one or more buses 826, one or more operating systems 828, and one or more computer-readable tangible storage devices 830.

[0051] Processor 820 is implemented in hardware, firmware, or a combination of hardware and software. Processor 820 is a central processing unit (CPU), a graphics processing unit (GPU), an accelerated processing unit (APU), a microprocessor, a microcontroller, a digital signal processor (DSP), a field-programmable gate array (FPGA), an application-specific integrated circuit (ASIC), or another type of processing component. In some implementations, processor 820 includes one or more processors capable of being programmed to perform a function. Bus 826 includes a component that permits communication among the internal components 800A,B.

[0052] The one or more operating systems 828, the software program 108 (FIG. 1) and the Point Cloud Compression Program 116 (FIG. 1) on server computer 114 (FIG. 1) are stored on one or more of the respective computer-readable tangible storage devices 830 for execution by one or more of the respective processors 820 via one or more of the respective RAMs 822 (which typically include cache memory). In the embodiment illustrated in FIG. 5, each of the computer-readable tangible storage devices 830 is a magnetic disk storage device of an internal hard drive. Alternatively, each of the computer-readable tangible storage devices 830 is a semiconductor

storage device such as ROM 824, EPROM, flash memory, an optical disk, a magneto-optic disk, a solid state disk, a compact disc (CD), a digital versatile disc (DVD), a floppy disk, a cartridge, a magnetic tape, and/or another type of non-transitory computer-readable tangible storage device that can store a computer program and digital information.

[0053] Each set of internal components 800A,B also includes a R/W drive or interface 832 to read from and write to one or more portable computer-readable tangible storage devices 936 such as a CD-ROM, DVD, memory stick, magnetic tape, magnetic disk, optical disk or semiconductor storage device. A software program, such as the software program 108 (FIG. 1) and the Point Cloud Compression Program 116 (FIG. 1) can be stored on one or more of the respective portable computer-readable tangible storage devices 936, read via the respective R/W drive or interface 832 and loaded into the respective hard drive 830.

[0054] Each set of internal components 800A,B also includes network adapters or interfaces 836 such as a TCP/IP adapter cards; wireless Wi-Fi interface cards; or 3G, 4G, or 5G wireless interface cards or other wired or wireless communication links. The software program 108 (FIG. 1) and the Point Cloud Compression Program 116 (FIG. 1) on the server computer 114 (FIG. 1) can be downloaded to the computer 102 (FIG. 1) and server computer 114 from an external computer via a network (for example, the Internet, a local area network or other, wide area network) and respective network adapters or interfaces 836. From the network adapters or interfaces 836, the software program 108 and the Point Cloud Compression Program 116 on the server computer 114 are loaded into the respective hard drive 830. The network may comprise copper wires, optical fibers, wireless transmission, routers, firewalls, switches, gateway computers and/or edge servers.

[0055] Each of the sets of external components 900A,B can include a computer display monitor 920, a keyboard 930, and a computer mouse 934. External components 900A,B can also include touch screens, virtual keyboards, touch pads, pointing devices, and other human interface devices. Each of the sets of internal components 800A,B also includes device drivers 840 to interface to computer display monitor 920, keyboard 930 and computer mouse 934. The device drivers 840, R/W drive or interface 832 and network adapter or interface 836 comprise hardware and software (stored in storage device 830 and/or ROM 824).

[0056] It is understood in advance that although this disclosure includes a detailed description

on cloud computing, implementation of the teachings recited herein are not limited to a cloud computing environment. Rather, some embodiments are capable of being implemented in conjunction with any other type of computing environment now known or later developed.

[0057] Cloud computing is a model of service delivery for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g. networks, network bandwidth, servers, processing, memory, storage, applications, virtual machines, and services) that can be rapidly provisioned and released with minimal management effort or interaction with a provider of the service. This cloud model may include at least five characteristics, at least three service models, and at least four deployment models.

[0058] Characteristics are as follows:

On-demand self-service: a cloud consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with the service's provider.

Broad network access: capabilities are available over a network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, laptops, and PDAs).

Resource pooling: the provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to demand. There is a sense of location independence in that the consumer generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at a higher level of abstraction (e.g., country, state, or datacenter).

Rapid elasticity: capabilities can be rapidly and elastically provisioned, in some cases automatically, to quickly scale out and rapidly released to quickly scale in. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be purchased in any quantity at any time.

Measured service: cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported providing transparency for both the provider and consumer of the utilized

service.

[0059] Service Models are as follows:

Software as a Service (SaaS): the capability provided to the consumer is to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through a thin client interface such as a web browser (e.g., web-based e-mail). The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.

Platform as a Service (PaaS): the capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including networks, servers, operating systems, or storage, but has control over the deployed applications and possibly application hosting environment configurations.

Infrastructure as a Service (IaaS): the capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, deployed applications, and possibly limited control of select networking components (e.g., host firewalls).

[0060] Deployment Models are as follows:

Private cloud: the cloud infrastructure is operated solely for an organization. It may be managed by the organization or a third party and may exist on-premises or off-premises.

Community cloud: the cloud infrastructure is shared by several organizations and supports a specific community that has shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be managed by the organizations or a third party and may exist on-premises or off-premises.

Public cloud: the cloud infrastructure is made available to the general public or a large industry group and is owned by an organization selling cloud services.

Hybrid cloud: the cloud infrastructure is a composition of two or more clouds (private,

community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load-balancing between clouds).

[0061] A cloud computing environment is service oriented with a focus on statelessness, low coupling, modularity, and semantic interoperability. At the heart of cloud computing is an infrastructure comprising a network of interconnected nodes.

[0062] Referring to FIG. 6, illustrative cloud computing environment 600 is depicted. As shown, cloud computing environment 600 comprises one or more cloud computing nodes 10 with which local computing devices used by cloud consumers, such as, for example, personal digital assistant (PDA) or cellular telephone 54A, desktop computer 54B, laptop computer 54C, and/or automobile computer system 54N may communicate. Cloud computing nodes 10 may communicate with one another. They may be grouped (not shown) physically or virtually, in one or more networks, such as Private, Community, Public, or Hybrid clouds as described hereinabove, or a combination thereof. This allows cloud computing environment 600 to offer infrastructure, platforms and/or software as services for which a cloud consumer does not need to maintain resources on a local computing device. It is understood that the types of computing devices 54A-N shown in FIG. 6 are intended to be illustrative only and that cloud computing nodes 10 and cloud computing environment 600 can communicate with any type of computerized device over any type of network and/or network addressable connection (e.g., using a web browser).

[0063] Referring to FIG. 7, a set of functional abstraction layers 700 provided by cloud computing environment 600 (FIG. 6) is shown. It should be understood in advance that the components, layers, and functions shown in FIG. 7 are intended to be illustrative only and embodiments are not limited thereto. As depicted, the following layers and corresponding functions are provided:

[0064] Hardware and software layer 60 includes hardware and software components. Examples of hardware components include: mainframes 61; RISC (Reduced Instruction Set Computer) architecture based servers 62; servers 63; blade servers 64; storage devices 65; and networks and networking components 66. In some embodiments, software components include network application server software 67 and database software 68.

[0065] Virtualization layer 70 provides an abstraction layer from which the following examples of virtual entities may be provided: virtual servers 71; virtual storage 72; virtual networks 73, including virtual private networks; virtual applications and operating systems 74; and virtual clients 75.

[0066] In one example, management layer 80 may provide the functions described below. Resource provisioning 81 provides dynamic procurement of computing resources and other resources that are utilized to perform tasks within the cloud computing environment. Metering and Pricing 82 provide cost tracking as resources are utilized within the cloud computing environment, and billing or invoicing for consumption of these resources. In one example, these resources may comprise application software licenses. Security provides identity verification for cloud consumers and tasks, as well as protection for data and other resources. User portal 83 provides access to the cloud computing environment for consumers and system administrators. Service level management 84 provides cloud computing resource allocation and management such that required service levels are met. Service Level Agreement (SLA) planning and fulfillment 85 provide pre-arrangement for, and procurement of, cloud computing resources for which a future requirement is anticipated in accordance with an SLA.

[0067] Workloads layer 90 provides examples of functionality for which the cloud computing environment may be utilized. Examples of workloads and functions which may be provided from this layer include: mapping and navigation 91; software development and lifecycle management 92; virtual classroom education delivery 93; data analytics processing 94; transaction processing 95; and Point Cloud Compression 96. Point Cloud Compression 96 may reduce an expanded set of neighboring nodes for a current node for compression and decompression of point cloud data.

[0068] Some embodiments may relate to a system, a method, and/or a computer readable medium at any possible technical detail level of integration. The computer readable medium may include a computer-readable non-transitory storage medium (or media) having computer readable program instructions thereon for causing a processor to carry out operations.

[0069] The computer readable storage medium can be a tangible device that can retain and store instructions for use by an instruction execution device. The computer readable storage medium may be, for example, but is not limited to, an electronic storage device, a magnetic storage

device, an optical storage device, an electromagnetic storage device, a semiconductor storage device, or any suitable combination of the foregoing. A non-exhaustive list of more specific examples of the computer readable storage medium includes the following: a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), a static random access memory (SRAM), a portable compact disc read-only memory (CD-ROM), a digital versatile disk (DVD), a memory stick, a floppy disk, a mechanically encoded device such as punch-cards or raised structures in a groove having instructions recorded thereon, and any suitable combination of the foregoing. A computer readable storage medium, as used herein, is not to be construed as being transitory signals per se, such as radio waves or other freely propagating electromagnetic waves, electromagnetic waves propagating through a waveguide or other transmission media (e.g., light pulses passing through a fiber-optic cable), or electrical signals transmitted through a wire.

[0070] Computer readable program instructions described herein can be downloaded to respective computing/processing devices from a computer readable storage medium or to an external computer or external storage device via a network, for example, the Internet, a local area network, a wide area network and/or a wireless network. The network may comprise copper transmission cables, optical transmission fibers, wireless transmission, routers, firewalls, switches, gateway computers and/or edge servers. A network adapter card or network interface in each computing/processing device receives computer readable program instructions from the network and forwards the computer readable program instructions for storage in a computer readable storage medium within the respective computing/processing device.

[0071] Computer readable program code/instructions for carrying out operations may be assembler instructions, instruction-set-architecture (ISA) instructions, machine instructions, machine dependent instructions, microcode, firmware instructions, state-setting data, configuration data for integrated circuitry, or either source code or object code written in any combination of one or more programming languages, including an object oriented programming language such as Smalltalk, C++, or the like, and procedural programming languages, such as the "C" programming language or similar programming languages. The computer readable program instructions may execute entirely on the user's computer, partly on the user's computer, as a stand-alone software package, partly on the user's computer and partly on a remote computer or entirely on the remote

computer or server. In the latter scenario, the remote computer may be connected to the user's computer through any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider). In some embodiments, electronic circuitry including, for example, programmable logic circuitry, field-programmable gate arrays (FPGA), or programmable logic arrays (PLA) may execute the computer readable program instructions by utilizing state information of the computer readable program instructions to personalize the electronic circuitry, in order to perform aspects or operations.

[0072] These computer readable program instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks. These computer readable program instructions may also be stored in a computer readable storage medium that can direct a computer, a programmable data processing apparatus, and/or other devices to function in a particular manner, such that the computer readable storage medium having instructions stored therein comprises an article of manufacture including instructions which implement aspects of the function/act specified in the flowchart and/or block diagram block or blocks.

[0073] The computer readable program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other device to cause a series of operational steps to be performed on the computer, other programmable apparatus or other device to produce a computer implemented process, such that the instructions which execute on the computer, other programmable apparatus, or other device implement the functions/acts specified in the flowchart and/or block diagram block or blocks.

[0074] The flowchart and block diagrams in the Figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods, and computer readable media according to various embodiments. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of instructions, which comprises one or more executable instructions for implementing the specified logical function(s). The method,

computer system, and computer readable medium may include additional blocks, fewer blocks, different blocks, or differently arranged blocks than those depicted in the Figures. In some alternative implementations, the functions noted in the blocks may occur out of the order noted in the Figures. For example, two blocks shown in succession may, in fact, be executed concurrently or substantially concurrently, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that perform the specified functions or acts or carry out combinations of special purpose hardware and computer instructions.

[0075] It will be apparent that systems and/or methods, described herein, may be implemented in different forms of hardware, firmware, or a combination of hardware and software. The actual specialized control hardware or software code used to implement these systems and/or methods is not limiting of the implementations. Thus, the operation and behavior of the systems and/or methods were described herein without reference to specific software code—it being understood that software and hardware may be designed to implement the systems and/or methods based on the description herein.

[0076] No element, act, or instruction used herein should be construed as critical or essential unless explicitly described as such. Also, as used herein, the articles “a” and “an” are intended to include one or more items, and may be used interchangeably with “one or more.” Furthermore, as used herein, the term “set” is intended to include one or more items (e.g., related items, unrelated items, a combination of related and unrelated items, etc.), and may be used interchangeably with “one or more.” Where only one item is intended, the term “one” or similar language is used. Also, as used herein, the terms “has,” “have,” “having,” or the like are intended to be open-ended terms. Further, the phrase “based on” is intended to mean “based, at least in part, on” unless explicitly stated otherwise.

[0077] The descriptions of the various aspects and embodiments have been presented for purposes of illustration, but are not intended to be exhaustive or limited to the embodiments disclosed. Even though combinations of features are recited in the claims and/or disclosed in the

specification, these combinations are not intended to limit the disclosure of possible implementations. In fact, many of these features may be combined in ways not specifically recited in the claims and/or disclosed in the specification. Although each dependent claim listed below may directly depend on only one claim, the disclosure of possible implementations includes each dependent claim in combination with every other claim in the claim set. Many modifications and variations will be apparent to those of ordinary skill in the art without departing from the scope of the described embodiments. The terminology used herein was chosen to best explain the principles of the embodiments, the practical application or technical improvement over technologies found in the marketplace, or to enable others of ordinary skill in the art to understand the embodiments disclosed herein.

CLAIMS

1. A method of decoding point cloud data, executable by a processor, comprising:
 - receiving data corresponding to a point cloud;
 - reducing a number of contexts associated with the received data based on occupancy data corresponding to one or more parent nodes and one or more child nodes within the received data;
 - decoding the data corresponding to the point cloud based on the reduced number of contexts;
 - wherein for parent nodes that share a same face with child nodes, a parent-level context index is calculated by taking all possible combinations, and for parent nodes that share an edge with child nodes, a parent-level context index is calculated by just counting how many of the parent nodes are occupied.
2. The method of claim 1, further comprising caching the occupancy data based on a hash table.
3. The method of claim 1, further comprising signaling a context reduction level in a sequence parameter set, a geometry parameter set, or a slice data header.
4. The method of claim 1, wherein the number of contexts is reduced based on one or more child nodes having a shortest distance from a current node.
5. The method of claim 1, wherein the number of contexts is reduced based on considering only a subset of child nodes from among the one or more child nodes.
6. The method of claim 1, wherein the number of contexts is reduced based on not differentiating between the one or more child nodes.
7. The method of claim 1, wherein the number of contexts is reduced based on discarding occupancy data associated with the one or more parent nodes.

8. A computer system for decoding point cloud data, the computer system comprising:
 - one or more computer-readable non-transitory storage media configured to store computer program code; and
 - one or more computer processors configured to access said computer program code and operate as instructed by said computer program code, said computer program code including:
 - receiving code configured to cause the one or more computer processors to receive data corresponding to a point cloud;
 - reducing code configured to cause the one or more computer processors to reduce a number of contexts associated with the received data based on occupancy data corresponding to one or more parent nodes and one or more child nodes within the received data;
 - decoding code configured to cause the one or more computer processors to decode the data corresponding to the point cloud based on the reduced number of contexts;
 - wherein for parent nodes that share a same face with child nodes, a parent-level context index is calculated by taking all possible combinations, and for parent nodes that share an edge with child nodes, a parent-level context index is calculated by just counting how many of the parent nodes are occupied.
9. The computer system of claim 8, further comprising caching code configured to cause the one or more computer processors to cache the occupancy data based on a hash table.
10. The computer system of claim 8, further comprising signaling code configured to cause the one or more computer processors to signal a context reduction level in a sequence parameter set, a geometry parameter set, or a slice data header.
11. The computer system of claim 8, wherein the number of contexts is reduced based on one or more child nodes having a shortest distance from a current node.
12. The computer system of claim 8, wherein the number of contexts is reduced based on considering only a subset of child nodes from among the one or more child nodes.

13. The computer system of claim 8, wherein the number of contexts is reduced based on not differentiating between the one or more child nodes.
14. The computer system of claim 8, wherein the number of contexts is reduced based on discarding occupancy data associated with the one or more parent nodes.
15. A non-transitory computer readable medium having stored thereon a computer program for decoding point cloud data, the computer program configured to cause one or more computer processors to:
 - receive data corresponding to a point cloud;
 - reduce a number of contexts associated with the received data based on occupancy data corresponding to one or more parent nodes and one or more child nodes within the received data;
 - decode the data corresponding to the point cloud based on the reduced number of contexts;
 - wherein for parent nodes that share a same face with child nodes, a parent-level context index is calculated by taking all possible combinations, and for parent nodes that share an edge with child nodes, a parent-level context index is calculated by just counting how many of the parent nodes are occupied.
16. The computer readable medium of claim 15, wherein the computer program is further configured to cause one or more computer processors to cache the occupancy data based on a hash table.
17. The computer readable medium of claim 15, wherein the computer program is further configured to cause one or more computer processors to signal a context reduction level in a sequence parameter set, a geometry parameter set, or a slice data header.
18. The computer readable medium of claim 15, wherein the number of contexts is reduced based on one or more child nodes having a shortest distance from a current node.

19. The computer readable medium of claim 15, wherein the number of contexts is reduced based on considering only a subset of child nodes from among the one or more child nodes.
20. The computer readable medium of claim 15, wherein the number of contexts is reduced based on not differentiating between the one or more child nodes.