

INTER PREDICTION CODING FOR GEOMETRY POINT CLOUD COMPRESSION

[0001] This application claims priority to U.S. Patent Application No. 18/303,390, filed 19 April 2023, U.S. Provisional Patent Application No. 63/363,306, filed 20 April 2022, and U.S. Provisional Patent Application No. 62/363,552, filed 25 April 2022, the entire content of each of which are incorporated herein by reference. U.S. Patent Application No. 18/303,390, filed 19 April 2023 claims the benefit of U.S. Provisional Patent Application No. 63/363,306, filed 20 April 2022, and U.S. Provisional Patent Application No. 62/363,552, filed 25 April 2022.

TECHNICAL FIELD

[0002] This disclosure relates to point cloud encoding and decoding.

BACKGROUND

[0003] A point cloud is a collection of points in a 3-dimensional space. The points may correspond to points on objects within the 3-dimensional space. Thus, a point cloud may be used to represent the physical content of the 3-dimensional space. Point clouds may have utility in a wide variety of situations. For example, point clouds may be used in the context of autonomous vehicles for representing the positions of objects on a roadway. In another example, point clouds may be used in the context of representing the physical content of an environment for purposes of positioning virtual objects in an augmented reality (AR) or mixed reality (MR) application. Point cloud compression is a process for encoding and decoding point clouds. Encoding point clouds may reduce the amount of data required for storage and transmission of point clouds.

SUMMARY

[0004] In general, this disclosure describes techniques for coding nodes of a point cloud using inter prediction, such as for the Geometry Point Cloud Compression (G-PCC) standard currently being developed. The example techniques, however, are not limited to the G-PCC standard. In some examples of G-PCC, coordinates of positions of nodes (also referred to as points) of a point cloud may be converted into a (r, ϕ, i) domain in which a position of a node is represented by three parameters, a radius r , an azimuth ϕ , and a laser index i (e.g., a laser identifier). When using an angular mode for predictive

geometry coding in G-PCC, a G-PCC coder may perform prediction in the (r, ϕ, i) domain. For instance, to code a particular node of a particular frame of a point cloud, the G-PCC coder may determine a predicted value for one or more of a radius r , azimuth ϕ , and laser index i of the particular node based on another node of the particular frame and add the predicted value to residual data (e.g., a residual radius r , a residual azimuth ϕ , or a residual laser index i) to determine a reconstructed radius r , azimuth ϕ , or laser index i of the particular node. As coding the residual data may account for most of the coding overhead, the coding efficiency (e.g., an amount of bits used to code a point) may be affected by how close the predicted value (e.g., the predicted radius r , azimuth ϕ , or laser index i) of the particular node is to the reconstructed value (e.g., the reconstructed radius r , azimuth ϕ , or laser index i) of the particular node. As such, generating more accurate predicted parameters for a node may improve overall coding performance.

[0005] In accordance with one or more techniques of this disclosure, a G-PCC coder may predict a current point of a current frame of a point cloud using inter prediction. For instance, to predict the current point using inter prediction, the G-PCC coder may identify a reference point in a different frame than the current frame and predict one or more parameters (e.g., a radius r , an azimuth ϕ , and a laser index i) of the current point based on one or more parameters of the reference point. In coding scenarios where the parameters of the reference point are closer to the parameters of the current point than the parent or other available points in the current frame, predicting the current point using inter prediction may reduce the size of residual data. In this way, the inter-prediction techniques of this disclosure may enable a G-PCC coder to improve coding efficiency.

[0006] According to an example of this disclosure, a device for processing (e.g., encoding or decoding) point cloud data includes: a memory configured to store the point cloud data; and one or more processors, implemented in circuitry, and configured to, responsive to a determination to predict a current point in a current frame using an inter prediction mode: identify a reference frame; derive one or more motion compensated reference points from the reference frame; selectively resample points in the reference frame based on the one or more motion compensated reference points to generate a resampled reference frame; and inter predict the current point in the current frame based on the resampled reference frame.

[0007] According to an example of this disclosure, a method of processing (e.g., encoding or decoding) a point cloud includes: responsive to determining to predict a current point in a current frame using an inter prediction mode: identifying a reference

frame; deriving one or more motion compensated reference points from the reference frame; selectively resampling points in the reference frame based on the one or more points motion compensated reference points to generate a resampled reference frame; and inter predicting the current point in the current frame based on the resampled reference frame.

[0008] A computer-readable storage medium stores instructions that when executed by one or more processors cause the one or more processors to: responsive to determining to predict a current point in a current frame using an inter prediction mode: identify a reference frame; derive one or more motion compensated reference points from the reference frame; selectively resample points in the reference frame based on the one or more points motion compensated reference points to generate a resampled reference frame; and inter predict the current point in the current frame based on the resampled reference frame.

[0009] The details of one or more examples are set forth in the accompanying drawings and the description below. Other features, objects, and advantages will be apparent from the description, drawings, and claims.

BRIEF DESCRIPTION OF DRAWINGS

[0010] FIG. 1 is a block diagram illustrating an example encoding and decoding system that may perform the techniques of this disclosure.

[0011] FIG. 2 is a block diagram illustrating an example Geometry Point Cloud Compression (G-PCC) encoder, in accordance with one or more aspects of this disclosure.

[0012] FIG. 3 is a block diagram illustrating an example G-PCC decoder, in accordance with one or more aspects of this disclosure.

[0013] FIG. 4 is a conceptual diagram illustrating octree split for geometry coding, in accordance with one or more aspects of this disclosure.

[0014] FIGS. 5A and 5B are conceptual diagrams of a spinning light detection and ranging (LIDAR) acquisition model, in accordance with one or more aspects of this disclosure.

[0015] FIG. 6 is a conceptual diagram of a prediction tree for predictive geometry coding, in accordance with one or more aspects of this disclosure.

[0016] FIG. 7 is a conceptual diagram illustrating an example inter prediction process for predicting points of a point cloud, in accordance with one or more aspects of this disclosure.

[0017] FIG. 8 is a flow diagram illustrating an example decoding flow associated with a syntax element indicating whether a node is coded in an inter prediction mode or an intra prediction mode.

[0018] FIG. 9 is a conceptual diagram illustrating an additional inter predictor point obtained from the first point that has azimuth greater than the inter predictor point.

[0019] FIG. 10 is a flowchart illustrating an example technique for how resampled points may be used for inter prediction, in accordance with one or more techniques of this disclosure.

[0020] FIG. 11 is a flowchart illustrating an example technique for how points in a global compensated reference frame may be resampled using points from a zero compensated reference frame, in accordance with one or more techniques of this disclosure.

[0021] FIGS. 12A-12D are conceptual diagrams illustrating how a first set of points in a global compensated reference frame may be chosen for a point in a zero compensated reference frame, in accordance with one or more aspects of this disclosure.

[0022] FIG. 13 is a flowchart illustrating a process in accordance with the techniques of this disclosure.

[0023] FIG. 14 is a conceptual diagram illustrating an example range-finding system that may be used with one or more techniques of this disclosure.

[0024] FIG. 15 is a conceptual diagram illustrating an example vehicle-based scenario in which one or more techniques of this disclosure may be used.

[0025] FIG. 16 is a conceptual diagram illustrating an example extended reality system in which one or more techniques of this disclosure may be used.

[0026] FIG. 17 is a conceptual diagram illustrating an example mobile device system in which one or more techniques of this disclosure may be used.

DETAILED DESCRIPTION

[0027] When using an angular mode for predictive geometry coding in G-PCC, a G-PCC coder may perform prediction in an angular domain, also referred to as the (r, ϕ, i) domain. For instance, to code a particular node of a particular frame of a point cloud, the G-PCC coder may determine a predicted value for one or more of a radius r , azimuth ϕ , or laser index i of the particular node based on another node of the particular frame and add the predicted radius r , azimuth ϕ , and laser index i of the node to residual data (e.g., a residual radius r , a residual azimuth ϕ , or a residual laser index i) to determine a reconstructed radius r , azimuth ϕ , or laser index i of the particular node. As coding the

residual data may account for significant coding overhead, the coding efficiency (e.g., an amount of bits used to code a point) may be affected by how close the predicted radius r , azimuth ϕ , and laser index i of the particular node is to the reconstructed radius r , azimuth ϕ , and laser index i of the particular node. As such, generating more accurate predicted parameters for a node may improve overall coding performance.

[0028] In accordance with one or more techniques of this disclosure, a G-PCC coder may predict a current point of a current frame of a point cloud using inter prediction. For instance, to predict the current point using inter prediction, the G-PCC coder may identify a reference point in a different frame than the current frame and predict one or more parameters (e.g., a radius r , an azimuth ϕ , and a laser index i) of the current point based on one or more parameters of the reference point. In coding scenarios where the parameters of the reference point are closer to the parameters of the current point than the parent or other available points in the current frame, predicting the current point using inter prediction may reduce the size of residual data compared to intra prediction. In this way, the inter-prediction techniques of this disclosure may enable a G-PCC coder to improve coding efficiency.

[0029] As will be explained in more detail below, inter prediction uses points in a reference frame to predict points in a current frame. In some examples, a zero-motion compensated, also referred to as zero compensated, reference frame may be used, and in some cases a global motion compensated reference frame may be used for prediction. In the zero compensated frame, the points in the reference frame are used without applying motion compensation. Accordingly, a zero-motion compensated reference frame may also at times in this disclosure simply be referred to as a reference frame. The global motion compensated reference frame may be obtained from the (zero compensated) reference frame by performing motion compensation based on a set of motion parameters. The motion parameters may be derived, for example, using motion estimation techniques or be obtained from external sources, such as using GPS parameters or from motion of an automobile capturing the point cloud. When motion compensation is applied to points in the reference frame, the locations of the points are modified, and this may results in a reduction in coding efficiency.

[0030] The reference frame contains points, and each point is associated with a reconstructed position. This reconstructed position is the position of the point with respect to a first position in the reference frame. For example, if the point was captured by a light detection and ranging (LIDAR) sensor mounted on a vehicle, the first position may

indicate a position associated with the vehicle, such as a position of a particular point in the vehicle/sensor. The position of the points in the reference frame may be relative to the first position. Compared to the reference frame, the vehicle may have moved and, therefore, a second position may be associated with the vehicle in the current frame such that positions of points in the current frame are relative to the second position. If there is motion between the reference frame and current frame, the first and second position are different. In this case, if motion compensation is applied, the position of one of more points in the reference frame (which were relative to the first position) may be modified such that positions of these points are relative to the second position. For example, suppose the vehicle was at location (1,1,1) in the reference frame and moved to location (3,4,5) in the second frame. The first position is (1,1,1), and the second position is (3,4,5). The motion parameter in this example is only translation, and the motion vector is (2,3,4) (obtained by subtracting 1,1,1 from 3,4,5, respectively). Suppose a point X in the reference frame has position (a, b, c) relative to the first position. After motion compensation, the position is modified to (a - 2, b - 3, c - 4). This may be considered to be the position of the point X with respect to the second position.

[0031] The above example only contains translation and was described in the cartesian domain. However motion parameters may also include rotation. For example, suppose, there are three successive points P1, P2, P3 that are captured by a laser of the LIDAR sensor in frame F1. As the sampling rate of the LIDAR sensor lasers are typically constant, the azimuth difference between successive points of a laser is the same (e.g., ΔAzim). Thus, the azimuth difference between P1 and P2 is the same (excluding noise) as the azimuth difference between P2 and P3, both being equal to ΔAzim . Consider another frame F2, where F1 is a reference frame for F2. When motion compensation is applied to points in F1, the position of the points P1, P2 and P3 are modified. Due to the operation of motion compensation (rotation, and in some cases translation), the difference in the azimuth values between compensated point of P1 and compensated point of P2 may not equal ΔAzim . Similarly, the difference in the azimuth values between compensated point of P2 and compensated point of P3 may not equal ΔAzim . When these compensated points are used for inter prediction of points in the current frame F2, where successive points of a laser have a azimuth difference of ΔAzim , this results in a sub-optimal predictor as the compensated points of P1, P2 and P3 may not correspond to any point in the current frame F2. This results in coding inefficiency.

[0032] This disclosure describes techniques to modify the compensated points of P1, P2 and P3 such that these points fall back on the regular sampling grid. This may be performed by looking at the azimuth values of the points in frame F1 (i.e., the zero compensated reference frame) and deriving new position values corresponding to these azimuth values from the compensated positions of points in frame F1. This results in a reference frame that has points that are closer to the sampling grid of the sensor, and therefore, a closer predictor to the points in the current frame. It must be noted that due to noise present, the point positions captured by a LIDAR sensor are typically not in a regular sampling grid; however, the sampling pattern and frequency are the same, or nearly the same, across frames in most cases.

[0033] According to the techniques of this disclosure, a G-PCC coder may, responsive to determining that a current point in a current frame is coded using an inter prediction mode, derive one or more motion compensated reference points from the reference frame and selectively resample points in the reference frame based on the one or more motion compensated reference points to generate a resampled reference frame. By inter predicting the current point in the current frame based on the resampled reference frame, a G-PCC coder may perform more accurate inter prediction, which reduces the bits needed to signal residual data, and thus improves compression without increasing distortion.

[0034] FIG. 1 is a block diagram illustrating an example encoding and decoding system 100 that may perform the techniques of this disclosure. The techniques of this disclosure are generally directed to coding (encoding and/or decoding) point cloud data, i.e., to support point cloud compression. In general, point cloud data includes any data for processing a point cloud. The coding may be effective in compressing and/or decompressing point cloud data.

[0035] As shown in FIG. 1, system 100 includes a source device 102 and a destination device 116. Source device 102 provides encoded point cloud data to be decoded by a destination device 116. Particularly, in the example of FIG. 1, source device 102 provides the point cloud data to destination device 116 via a computer-readable medium 110. Source device 102 and destination device 116 may comprise any of a wide range of devices, including desktop computers, notebook (i.e., laptop) computers, tablet computers, set-top boxes, telephone handsets such as smartphones, televisions, cameras, display devices, digital media players, video gaming consoles, video streaming devices, terrestrial or marine vehicles, spacecraft, aircraft, robots, LIDAR devices, satellites, or

the like. In some cases, source device 102 and destination device 116 may be equipped for wireless communication.

[0036] In the example of FIG. 1, source device 102 includes a data source 104, a memory 106, a G-PCC encoder 200, and an output interface 108. Destination device 116 includes an input interface 122, a G-PCC decoder 300, a memory 120, and a data consumer 118. In accordance with this disclosure, G-PCC encoder 200 of source device 102 and G-PCC decoder 300 of destination device 116 may be configured to apply the techniques of this disclosure related to predictive geometry coding. Thus, source device 102 represents an example of an encoding device, while destination device 116 represents an example of a decoding device. In other examples, source device 102 and destination device 116 may include other components or arrangements. For example, source device 102 may receive data (e.g., point cloud data) from an internal or external source. Likewise, destination device 116 may interface with an external data consumer, rather than include a data consumer in the same device.

[0037] System 100 as shown in FIG. 1 is merely one example. In general, other digital encoding and/or decoding devices may perform the techniques of this disclosure related to predictive geometry coding. Source device 102 and destination device 116 are merely examples of such devices in which source device 102 generates coded data for transmission to destination device 116. This disclosure refers to a “coding” device as a device that performs coding (encoding and/or decoding) of data. Thus, G-PCC encoder 200 and G-PCC decoder 300 represent examples of coding devices, in particular, an encoder and a decoder, respectively. In some examples, source device 102 and destination device 116 may operate in a substantially symmetrical manner such that each of source device 102 and destination device 116 includes encoding and decoding components. Hence, system 100 may support one-way or two-way transmission between source device 102 and destination device 116, e.g., for streaming, playback, broadcasting, telephony, navigation, and other applications.

[0038] In general, data source 104 represents a source of data (i.e., raw, unencoded point cloud data) and may provide a sequential series of “frames” of the data to G-PCC encoder 200, which encodes data for the frames. Data source 104 of source device 102 may include a point cloud capture device, such as any of a variety of cameras or sensors, e.g., a 3D scanner or a LIDAR device, one or more video cameras, an archive containing previously captured data, and/or a data feed interface to receive data from a data content provider. Alternatively or additionally, point cloud data may be computer-generated from

scanner, camera, sensor or other data. For example, data source 104 may generate computer graphics-based data as the source data, or produce a combination of live data, archived data, and computer-generated data. In each case, G-PCC encoder 200 encodes the captured, pre-captured, or computer-generated data. G-PCC encoder 200 may rearrange the frames from the received order (sometimes referred to as “display order”) into a coding order for coding. G-PCC encoder 200 may generate one or more bitstreams including encoded data. Source device 102 may then output the encoded data via output interface 108 onto computer-readable medium 110 for reception and/or retrieval by, e.g., input interface 122 of destination device 116.

[0039] Memory 106 of source device 102 and memory 120 of destination device 116 may represent general purpose memories. In some examples, memory 106 and memory 120 may store raw data, e.g., raw data from data source 104 and raw, decoded data from G-PCC decoder 300. Additionally or alternatively, memory 106 and memory 120 may store software instructions executable by, e.g., G-PCC encoder 200 and G-PCC decoder 300, respectively. Although memory 106 and memory 120 are shown separately from G-PCC encoder 200 and G-PCC decoder 300 in this example, it should be understood that G-PCC encoder 200 and G-PCC decoder 300 may also include internal memories for functionally similar or equivalent purposes. Furthermore, memory 106 and memory 120 may store encoded data, e.g., output from G-PCC encoder 200 and input to G-PCC decoder 300. In some examples, portions of memory 106 and memory 120 may be allocated as one or more buffers, e.g., to store raw, decoded, and/or encoded data. For instance, memory 106 and memory 120 may store data representing a point cloud.

[0040] Computer-readable medium 110 may represent any type of medium or device capable of transporting the encoded data from source device 102 to destination device 116. In one example, computer-readable medium 110 represents a communication medium to enable source device 102 to transmit encoded data directly to destination device 116 in real-time, e.g., via a radio frequency network or computer-based network. Output interface 108 may modulate a transmission signal including the encoded data, and input interface 122 may demodulate the received transmission signal, according to a communication standard, such as a wireless communication protocol. The communication medium may comprise any wireless or wired communication medium, such as a radio frequency (RF) spectrum or one or more physical transmission lines. The communication medium may form part of a packet-based network, such as a local area network, a wide-area network, or a global network such as the Internet. The

communication medium may include routers, switches, base stations, or any other equipment that may be useful to facilitate communication from source device 102 to destination device 116.

[0041] In some examples, source device 102 may output encoded data from output interface 108 to storage device 112. Similarly, destination device 116 may access encoded data from storage device 112 via input interface 122. Storage device 112 may include any of a variety of distributed or locally accessed data storage media such as a hard drive, Blu-ray discs, DVDs, CD-ROMs, flash memory, volatile or non-volatile memory, or any other suitable digital storage media for storing encoded data.

[0042] In some examples, source device 102 may output encoded data to file server 114 or another intermediate storage device that may store the encoded data generated by source device 102. Destination device 116 may access stored data from file server 114 via streaming or download. File server 114 may be any type of server device capable of storing encoded data and transmitting that encoded data to the destination device 116. File server 114 may represent a web server (e.g., for a website), a File Transfer Protocol (FTP) server, a content delivery network device, or a network attached storage (NAS) device. Destination device 116 may access encoded data from file server 114 through any standard data connection, including an Internet connection. This may include a wireless channel (e.g., a Wi-Fi connection), a wired connection (e.g., digital subscriber line (DSL), cable modem, etc.), or a combination of both that is suitable for accessing encoded data stored on file server 114. File server 114 and input interface 122 may be configured to operate according to a streaming transmission protocol, a download transmission protocol, or a combination thereof.

[0043] Output interface 108 and input interface 122 may represent wireless transmitters/receivers, modems, wired networking components (e.g., Ethernet cards), wireless communication components that operate according to any of a variety of IEEE 802.11 standards, or other physical components. In examples where output interface 108 and input interface 122 comprise wireless components, output interface 108 and input interface 122 may be configured to transfer data, such as encoded data, according to a cellular communication standard, such as 4G, 4G-LTE (Long-Term Evolution), LTE Advanced, 5G, or the like. In some examples where output interface 108 comprises a wireless transmitter, output interface 108 and input interface 122 may be configured to transfer data, such as encoded data, according to other wireless standards, such as an IEEE 802.11 specification, an IEEE 802.15 specification (e.g., ZigBee™), a Bluetooth™

standard, or the like. In some examples, source device 102 and/or destination device 116 may include respective system-on-a-chip (SoC) devices. For example, source device 102 may include an SoC device to perform the functionality attributed to G-PCC encoder 200 and/or output interface 108, and destination device 116 may include an SoC device to perform the functionality attributed to G-PCC decoder 300 and/or input interface 122.

[0044] The techniques of this disclosure may be applied to encoding and decoding in support of any of a variety of applications, such as communication between autonomous vehicles, communication between scanners, cameras, sensors and processing devices such as local or remote servers, geographic mapping, or other applications.

[0045] Input interface 122 of destination device 116 receives an encoded bitstream from computer-readable medium 110 (e.g., a communication medium, storage device 112, file server 114, or the like). The encoded bitstream may include signaling information defined by G-PCC encoder 200, which is also used by G-PCC decoder 300, such as syntax elements having values that describe characteristics and/or processing of coded units (e.g., slices, pictures, groups of pictures, sequences, or the like). Data consumer 118 uses the decoded data. For example, data consumer 118 may use the decoded data to determine the locations of physical objects. In some examples, data consumer 118 may comprise a display to present imagery based on a point cloud.

[0046] G-PCC encoder 200 and G-PCC decoder 300 each may be implemented as any of a variety of suitable encoder and/or decoder circuitry, such as one or more microprocessors, digital signal processors (DSPs), application specific integrated circuits (ASICs), field programmable gate arrays (FPGAs), discrete logic, software, hardware, firmware or any combinations thereof. When the techniques are implemented partially in software, a device may store instructions for the software in a suitable, non-transitory computer-readable medium and execute the instructions in hardware using one or more processors to perform the techniques of this disclosure. Each of G-PCC encoder 200 and G-PCC decoder 300 may be included in one or more encoders or decoders, either of which may be integrated as part of a combined encoder/decoder (CODEC) in a respective device. A device including G-PCC encoder 200 and/or G-PCC decoder 300 may comprise one or more integrated circuits, microprocessors, and/or other types of devices.

[0047] G-PCC encoder 200 and G-PCC decoder 300 may operate according to a coding standard, such as video point cloud compression (V-PCC) standard or a geometry point cloud compression (G-PCC) standard. This disclosure may generally refer to coding (e.g., encoding and decoding) of pictures to include the process of encoding or decoding

data. An encoded bitstream generally includes a series of values for syntax elements representative of coding decisions (e.g., coding modes).

[0048] This disclosure may generally refer to “signaling” certain information, such as syntax elements. The term “signaling” may generally refer to the communication of values for syntax elements and/or other data used to decode encoded data. That is, G-PCC encoder 200 may signal values for syntax elements in the bitstream. In general, signaling refers to generating a value in the bitstream. As noted above, source device 102 may transport the bitstream to destination device 116 substantially in real time, or not in real time, such as might occur when storing syntax elements to storage device 112 for later retrieval by destination device 116.

[0049] ISO/IEC MPEG (JTC 1/SC 29/WG 11) is studying the potential need for standardization of point cloud coding technology with a compression capability that significantly exceeds that of the current approaches and will target to create the standard. The group is working together on this exploration activity in a collaborative effort known as the 3-Dimensional Graphics Team (3DG) to evaluate compression technology designs proposed by their experts in this area.

[0050] Point cloud compression activities are categorized in two different approaches. The first approach is “Video point cloud compression” (V-PCC), which segments the 3D object, and project the segments in multiple 2D planes (which are represented as “patches” in the 2D frame), which are further coded by a legacy 2D video codec such as a High Efficiency Video Coding (HEVC) (ITU-T H.265) codec. The second approach is “Geometry-based point cloud compression” (G-PCC), which directly compresses 3D geometry i.e., position of a set of points in 3D space, and associated attribute values (for each point associated with the 3D geometry). G-PCC addresses the compression of point clouds in both Category 1 (static point clouds) and Category 3 (dynamically acquired point clouds). A draft of the G-PCC standard is available in G-PCC DIS, ISO/IEC JTC1/SC29/WG11 w19328, Brussels, Belgium, January 2020, and a description of the codec is available in G-PCC Codec Description v8, ISO/IEC JTC1/SC29/WG11 w19525, Brussels, Belgium, January 2020.

[0051] A point cloud contains a set of points in a 3D space, and may have attributes associated with the point. The attributes may be color information such as R, G, B or Y, Cb, Cr, or reflectance information, or other attributes. Point clouds may be captured by a variety of cameras or sensors such as LIDAR sensors and 3D scanners and may also be computer-generated. Point cloud data are used in a variety of applications including, but

not limited to, construction (modeling), graphics (3D models for visualizing and animation), and the automotive industry (LIDAR sensors used to help in navigation).

[0052] The 3D space occupied by a point cloud data may be enclosed by a virtual bounding box. The position of the points in the bounding box may be represented by a certain precision; therefore, the positions of one or more points may be quantized based on the precision. At the smallest level, the bounding box is split into voxels which are the smallest unit of space represented by a unit cube. A voxel in the bounding box may be associated with zero, one, or more than one point. The bounding box may be split into multiple cube/cuboid regions, which may be called tiles. Each tile may be coded into one or more slices. The partitioning of the bounding box into slices and tiles may be based on number of points in each partition, or based on other considerations (e.g., a particular region may be coded as tiles). The slice regions may be further partitioned using splitting decisions similar to those in video codecs.

[0053] FIG. 2 provides an overview of G-PCC encoder 200. FIG. 3 provides an overview of G-PCC decoder 300. The modules shown are logical, and do not necessarily correspond one-to-one to implemented code in the reference implementation of G-PCC codec, i.e., TMC13 test model software studied by ISO/IEC MPEG (JTC 1/SC 29/WG 11).

[0054] In both G-PCC encoder 200 and G-PCC decoder 300, point cloud positions are coded first. Attribute coding depends on the decoded geometry. In FIG. 2 and FIG. 3, modules 212, 218, 310, and 314 are options typically used for Category 1 data. Modules 220, 222, 316, and 318 are options typically used for Category 3 data. All the other modules are common between Categories 1 and 3.

[0055] For Category 3 data, the compressed geometry is typically represented as an octree from the root all the way down to a leaf level of individual voxels. For Category 1 data, the compressed geometry is typically represented by a pruned octree (i.e., an octree from the root down to a leaf level of blocks larger than voxels) plus a model that approximates the surface within each leaf of the pruned octree. In this way, both Category 1 and 3 data share the octree coding mechanism, while Category 1 data may in addition approximate the voxels within each leaf with a surface model. The surface model used is a triangulation comprising 1-10 triangles per block, resulting in a triangle soup. The Category 1 geometry codec is therefore known as the Trisoup geometry codec, while the Category 3 geometry codec is known as the Octree geometry codec.

[0056] FIG. 4 is a conceptual diagram illustrating an example octree split for geometry coding according to the techniques of this disclosure. In the example shown in FIG. 4, octree 400, may be split into a series of nodes. For example, each node may be a cubic node. At each node of an octree, G-PCC encoder 200 may signal an occupancy of a node by a point of the point cloud to G-PCC decoder 300, when the occupancy is not inferred by G-PCC decoder 300, for one or more of the node's child nodes, which may include up to eight nodes. Multiple neighborhoods are specified including (a) nodes that share a face with a current octree node, (b) nodes that share a face, edge, or a vertex with the current octree node, etc. Within each neighborhood, the occupancy of a node and/or its children may be used to predict the occupancy of the current node or its children. For points that are sparsely populated in certain nodes of the octree, the codec also supports a direct coding mode where the 3D position of the point is encoded directly. A flag may be signaled to indicate that a direct mode is signaled. With a direct mode, positions of points in the point cloud may be coded directly without any compression. At the lowest level, the number of points associated with the octree node/leaf node may also be coded.

[0057] Once the geometry is coded, the attributes corresponding to the geometry points are coded. When there are multiple attribute points corresponding to one reconstructed/decoded geometry point, an attribute value may be derived that is representative of the reconstructed point.

[0058] There are three attribute coding processes in G-PCC: Region Adaptive Hierarchical Transform (RAHT) coding, interpolation-based hierarchical nearest-neighbor prediction (Predicting Transform), and interpolation-based hierarchical nearest-neighbor prediction with an update/lifting step (Lifting Transform). RAHT and Lifting Transform are typically used for Category 1 data, while Predicting Transform is typically used for Category 3 data. However, either process may be used for any data, and, just like with the geometry codecs in G-PCC, the attribute coding process used to code the point cloud is specified in the bitstream.

[0059] The coding of the attributes may be conducted in a level-of-detail (LOD), where with each level of detail a finer representation of the point cloud attribute may be obtained. Each level of detail may be specified based on a distance metric from the neighboring nodes or based on a sampling distance.

[0060] At G-PCC encoder 200, the residuals obtained as the output of the coding processes for the attributes are quantized (e.g., by one of arithmetic encoding units 214

and/or 226). The quantized residuals may be coded using context adaptive arithmetic coding.

[0061] In the example of FIG. 2, G-PCC encoder 200 may include a coordinate transform unit 202, a color transform unit 204, a voxelization unit 206, an attribute transfer unit 208, an octree analysis unit 210, a surface approximation analysis unit 212, an arithmetic encoding unit 214, a geometry reconstruction unit (GRU) 216, an RAHT unit 218, a LOD generation unit 220, a lifting unit 222, a coefficient quantization unit 224, and an arithmetic encoding unit 226.

[0062] As shown in the example of FIG. 2, G-PCC encoder 200 may receive a set of positions and a set of attributes. The positions may include coordinates of points in a point cloud. The attributes may include information about points in the point cloud, such as colors associated with points in the point cloud.

[0063] Coordinate transform unit 202 may apply a transform to the coordinates of the points to transform the coordinates from an initial domain to a transform domain. This disclosure may refer to the transformed coordinates as transform coordinates. Color transform unit 204 may apply a transform to transform color information of the attributes to a different domain. For example, color transform unit 204 may transform color information from an RGB color space to a YCbCr color space.

[0064] Furthermore, in the example of FIG. 2, voxelization unit 206 may voxelize the transform coordinates. Voxelization of the transform coordinates may include quantization and removing some points of the point cloud. In other words, multiple points of the point cloud may be subsumed within a single “voxel,” which may thereafter be treated in some respects as one point. Furthermore, octree analysis unit 210 may generate an octree based on the voxelized transform coordinates. Additionally, in the example of FIG. 2, surface approximation analysis unit 212 may analyze the points to potentially determine a surface representation of sets of the points. Arithmetic encoding unit 214 may entropy encode syntax elements representing the information of the octree and/or surfaces determined by surface approximation analysis unit 212. G-PCC encoder 200 may output these syntax elements in a geometry bitstream.

[0065] Geometry reconstruction unit 216 may reconstruct transform coordinates of points in the point cloud based on the octree, data indicating the surfaces determined by surface approximation analysis unit 212, and/or other information. The number of transform coordinates reconstructed by geometry reconstruction unit 216 may be different from the original number of points of the point cloud because of voxelization and surface

approximation. This disclosure may refer to the resulting points as reconstructed points. Attribute transfer unit 208 may transfer attributes of the original points of the point cloud to reconstructed points of the point cloud. As shown in FIG. 2, attribute transfer unit 208 may transfer the attributes to one or both of RAHT unit 218 and LOD generation unit 220.

[0066] Furthermore, RAHT unit 218 may apply RAHT coding to the attributes of the reconstructed points. Alternatively or additionally, LOD generation unit 220 and lifting unit 222 may apply LOD processing and lifting, respectively, to the attributes of the reconstructed points. RAHT unit 218 and lifting unit 222 may generate coefficients based on the attributes. Coefficient quantization unit 224 may quantize the coefficients generated by RAHT unit 218 or lifting unit 222. Arithmetic encoding unit 226 may apply arithmetic coding to syntax elements representing the quantized coefficients. G-PCC encoder 200 may output these syntax elements in an attribute bitstream.

[0067] In the example of FIG. 3, G-PCC decoder 300 may include a geometry arithmetic decoding unit 302, an attribute arithmetic decoding unit 304, an octree synthesis unit 306, an inverse quantization unit 308, a surface approximation synthesis unit 310, a geometry reconstruction unit 312, a RAHT unit 314, a LoD generation unit 316, an inverse lifting unit 318, an inverse transform coordinate unit 320, and an inverse transform color unit 322.

[0068] G-PCC decoder 300 may obtain a geometry bitstream and an attribute bitstream. Geometry arithmetic decoding unit 302 of decoder 300 may apply arithmetic decoding (e.g., Context-Adaptive Binary Arithmetic Coding (CABAC) or other type of arithmetic decoding) to syntax elements in the geometry bitstream. Similarly, attribute arithmetic decoding unit 304 may apply arithmetic decoding to syntax elements in the attribute bitstream.

[0069] Octree synthesis unit 306 may synthesize an octree based on syntax elements parsed from the geometry bitstream. In instances where surface approximation is used in the geometry bitstream, surface approximation synthesis unit 310 may determine a surface model based on syntax elements parsed from the geometry bitstream and based on the octree.

[0070] Furthermore, geometry reconstruction unit 312 may perform a reconstruction to determine coordinates of points in a point cloud. Inverse transform coordinate unit 320 may apply an inverse transform to the reconstructed coordinates to convert the

reconstructed coordinates (positions) of the points in the point cloud from a transform domain back into an initial domain.

[0071] Additionally, in the example of FIG. 3, inverse quantization unit 308 may inverse quantize attribute values. The attribute values may be based on syntax elements obtained from the attribute bitstream (e.g., including syntax elements decoded by attribute arithmetic decoding unit 304). As shown in FIG. 3, inverse quantization unit 308 may transfer the attribute values to one or both of RAHT unit 314 and LOD generation unit 316.

[0072] Depending on how the attribute values are encoded, RAHT unit 314 may perform RAHT coding to determine, based on the inverse quantized attribute values, color values for points of the point cloud. Alternatively, LOD generation unit 316 and inverse lifting unit 318 may determine color values for points of the point cloud using a level of detail-based technique.

[0073] Furthermore, in the example of FIG. 3, inverse transform color unit 322 may apply an inverse color transform to the color values. The inverse color transform may be an inverse of a color transform applied by color transform unit 204 of G-PCC encoder 200. For example, color transform unit 204 may transform color information from an RGB color space to a YCbCr color space. Accordingly, inverse color transform unit 322 may transform color information from the YCbCr color space to the RGB color space.

[0074] The various units of FIG. 2 and FIG. 3 are illustrated to assist with understanding the operations performed by G-PCC encoder 200 and G-PCC decoder 300. The units may be implemented as fixed-function circuits, programmable circuits, or a combination thereof. Fixed-function circuits refer to circuits that provide particular functionality, and are preset on the operations that can be performed. Programmable circuits refer to circuits that can be programmed to perform various tasks, and provide flexible functionality in the operations that can be performed. For instance, programmable circuits may execute software or firmware that cause the programmable circuits to operate in the manner defined by instructions of the software or firmware. Fixed-function circuits may execute software instructions (e.g., to receive parameters or output parameters), but the types of operations that the fixed-function circuits perform are generally immutable. In some examples, one or more of the units may be distinct circuit blocks (fixed-function or programmable), and in some examples, one or more of the units may be integrated circuits.

[0075] Predictive geometry coding was introduced as an alternative to the octree geometry coding, where the nodes are arranged in a tree structure (which defines the prediction structure), and various prediction strategies are used to predict the coordinates of each node in the tree with respect to its predictors. FIG. 6 is a conceptual diagram illustrating an example of a prediction tree 600, a directed graph where the arrow points to the prediction direction. The horizontally shaded node is the root vertex and has no predictors; the grid shaded nodes have two children; the diagonally shaded node has 3 children; the non-shaded nodes have one children and the vertically shaded nodes are leaf nodes and these have no children. Every node has only one parent node.

[0076] Four prediction strategies may be specified for each node based on its parent (p0), grand-parent (p1) and great-grand-parent (p2). The prediction strategies include, no prediction, delta prediction (p0), linear prediction ($2 \cdot p0 - p1$), and parallelogram prediction ($p0 + p1 - p2$).

[0077] The encoder (e.g., G-PCC encoder 200) may employ any algorithm to generate the prediction tree; the algorithm used may be determined based on the application/use case and several strategies may be used. The encoder may encode, for each node, the residual coordinate values in the bitstream starting from the root node in a depth-first manner. Predictive geometry coding may be particularly useful for Category 3 (LIDAR-acquired) point cloud data e.g., for low-latency applications.

[0078] Angular mode may be used in predictive geometry coding, where the characteristics of LIDAR sensors may be utilized in coding the prediction tree more efficiently. The coordinates of the positions are converted to the (r, ϕ, i) (radius, azimuth, and laser index) and a prediction is performed in this domain (the residuals are coded in r, ϕ, i domain). Due to the errors in rounding, coding in r, ϕ, i is not lossless and hence a second set of residuals may be coded which correspond to the Cartesian coordinates. A description of the encoding and decoding strategies used for angular mode for predictive geometry coding is reproduced below. The description is based on FIGS. 5A and 5B, which are conceptual diagrams of a spinning LIDAR acquisition model.

[0079] The techniques of this disclosure may be applicable to at least point clouds acquired using a spinning Lidar model. Here, the lidar 502 has N lasers (e.g., N=16, 32, 64) spinning around the Z axis according to an azimuth angle ϕ (see FIGS. 5A and 5B). Each laser may have different elevation $\theta(i)_{i=1 \dots N}$ and height $\zeta(i)_{i=1 \dots N}$. The laser i hits a point M , with cartesian integer coordinates (x, y, z) , defined according to coordinate system 500 described in FIG. 5A.

[0080] The position of M is modelled with three parameters (r, ϕ, i) , which may be computed as follows:

- $r = \sqrt{x^2 + y^2}$
- $\phi = \text{atan2}(y, x)$
- $i = \arg \min_{j=1 \dots N} \{z + \zeta(j) - r \times \tan(\theta(j))\}$,

[0081] The coding process may use the quantized version of (r, ϕ, i) , denoted $(\tilde{r}, \tilde{\phi}, i)$, where the three integers $\tilde{r}, \tilde{\phi}$ and i may be computed as follows:

- $\tilde{r} = \text{floor}\left(\frac{\sqrt{x^2 + y^2}}{q_r} + o_r\right) = \text{hypot}(x, y)$
- $\tilde{\phi} = \text{sign}(\text{atan2}(y, x)) \times \text{floor}\left(\frac{|\text{atan2}(y, x)|}{q_\phi} + o_\phi\right)$
- $i = \arg \min_{j=1 \dots N} \{z + \zeta(j) - r \times \tan(\theta(j))\}$

where

- (q_r, o_r) and (q_ϕ, o_ϕ) are quantization parameters controlling the precision of $\tilde{\phi}$ and $\tilde{r}$, respectively.
- $\text{sign}(t)$ is the function that returns 1 if t is positive and (-1) otherwise.
- $|t|$ is the absolute value of t .

[0082] To avoid reconstruction mismatches due to the use of floating-point operations, the values of $\zeta(i)_{i=1 \dots N}$ and $\tan(\theta(i))_{i=1 \dots N}$ may be pre-computed and quantized as follows:

$$\tilde{z}(i) = \text{sign}(\zeta(i)) \times \text{floor}\left(\frac{|\zeta(i)|}{q_\zeta} + o_\zeta\right)$$

$$\tilde{t}(i) = \text{sign}(\zeta(\tan(\theta(j)))) \times \text{floor}\left(\frac{|\tan(\theta(j))|}{q_\theta} + o_\theta\right)$$

where

- (q_ζ, o_ζ) and (q_θ, o_θ) are quantization parameters controlling the precision of $\tilde{\zeta}$ and $\tilde{\theta}$, respectively.

The reconstructed cartesian coordinates are obtained as follows:

- $\hat{x} = \text{round}(\tilde{r} \times q_r \times \text{app_cos}(\tilde{\phi} \times q_\phi))$
- $\hat{y} = \text{round}(\tilde{r} \times q_r \times \text{app_sin}(\tilde{\phi} \times q_\phi))$
- $\hat{z} = \text{round}(\tilde{r} \times q_r \times \tilde{t}(i) \times q_\theta - \tilde{z}(i) \times q_\zeta)$,

where $\text{app_cos}(\cdot)$ and $\text{app_sin}(\cdot)$ are approximation of $\cos(\cdot)$ and $\sin(\cdot)$. The calculations could be using a fixed-point representation, a look-up table and linear interpolation.

[0083] In some examples, $(\hat{x}, \hat{y}, \hat{z})$ may be different from (x, y, z) due to various reasons, including:

- quantization
- approximations
- model imprecision
- model parameters imprecisions

[0084] In some examples, the reconstruction residuals (r_x, r_y, r_z) can be defined as follows:

- $r_x = x - \hat{x}$
- $r_y = y - \hat{y}$
- $r_z = z - \hat{z}$

[0085] In this process, G-PCC encoder 200 may proceed as follows:

- Encode the model parameters $\tilde{t}(i)$ and $\tilde{z}(i)$ and the quantization parameters q_ζ , q_θ and q_ϕ
- Apply the geometry predictive scheme described in G-PCC DIS to the representation $(\tilde{r}, \tilde{\phi}, i)$
 - A new predictor leveraging the characteristics of lidar could be introduced. For instance, the rotation speed of the lidar scanner around the z-axis is usually constant. Therefore, the G-PCC decoder may predict the current $\tilde{\phi}(j)$ as follows:

$$\tilde{\phi}(j) = \tilde{\phi}(j-1) + n(j) \times \delta_\phi(k)$$

Where

- $(\delta_\phi(k))_{k=1\dots K}$ is a set of potential speeds the encoder could choose from. The index k could be explicitly written to the bitstream or could be inferred from the context based on a deterministic strategy applied by both the encoder and the decoder, and
- $n(j)$ is the number of skipped points, which could be explicitly written to the bitstream or could be inferred from the context based on a deterministic strategy applied by both the encoder and the decoder.
- Encode with each node the reconstruction residuals (r_x, r_y, r_z)

[0086] G-PCC decoder 300 may proceed as follows:

- Decode the model parameters $\tilde{t}(i)$ and $\tilde{z}(i)$ and the quantization parameters q_r , q_ζ , q_θ and q_ϕ
- Decode the $(\tilde{r}, \tilde{\phi}, i)$ parameters associated with the nodes according to the geometry predictive scheme described in G-PCC Draft International Standard (DIS)
- Compute the reconstructed coordinates $(\hat{x}, \hat{y}, \hat{z})$ as described above
- Decode the residuals (r_x, r_y, r_z)
 - As discussed in the next section, lossy compression could be supported by quantizing the reconstruction residuals (r_x, r_y, r_z)
- Compute the original coordinates (x, y, z) as follows
 - $x = r_x + \hat{x}$
 - $y = r_y + \hat{y}$
 - $z = r_z + \hat{z}$

[0087] Lossy compression could be achieved by applying quantization to the reconstruction residuals (r_x, r_y, r_z) or by dropping points.

[0088] The quantized reconstruction residuals are computed as follows:

- $\tilde{r}_x = \text{sign}(r_x) \times \text{floor}\left(\frac{|r_x|}{q_x} + o_x\right)$
- $\tilde{r}_y = \text{sign}(r_y) \times \text{floor}\left(\frac{|r_y|}{q_y} + o_y\right)$
- $\tilde{r}_z = \text{sign}(r_z) \times \text{floor}\left(\frac{|r_z|}{q_z} + o_z\right)$

Where (q_x, o_x) , (q_y, o_y) and (q_z, o_z) are quantization parameters controlling the precision of $\tilde{r}_x$, $\tilde{r}_y$ and $\tilde{r}_z$, respectively.

[0089] In some examples, G-PCC encoder 200 and/or G-PCC decoder 300 may use Trellis quantization to further improve the RD (rate-distortion) performance results. The quantization parameters may change at sequence/frame/slice/block level to achieve region adaptive quality and for rate control purposes.

[0090] A G-PCC coder (e.g., G-PCC encoder 200 and/or G-PCC decoder 300) may perform point cloud compression using inter-prediction. By using inter-prediction, the G-PCC coder may use redundancy of points across frames to provide additional bit rate savings. Examples in the various aspects of this disclosure may be used individually or in any combination.

[0091] Although the discussion is predominantly on the polar coordinate system, processes disclosed in this application may also apply to other coordinate system such as

Cartesian, spherical, or any custom coordinate system that may be used to represent/code the point cloud positions and attributes.

[0092] The G-PCC coder may determine whether to code a point using inter-prediction or intra prediction. For instance, the G-PCC encoder may perform an analysis to determine it whether would be beneficial (e.g., in terms of bitrate or other conditions) to code a particular point using inter or intra prediction. In some examples, the G-PCC coder may perform this determination in one of the following ways. In some examples, G-PCC encoder 200 may use a first set of conditions during encoding, and G-PCC decoder 300 may use the first set of conditions during decoding. In some examples, G-PCC encoder 200 may use a first set of conditions during encoding and signal a syntax element that indicates the inter-prediction mode, and G-PCC decoder 300 may determine the mode based on the signaled syntax element.

[0093] The first set of conditions may be derived using one or more characteristics of the point such as cartesian position, angular coordinates (radius, azimuth, elevation), prediction mode of neighboring nodes, etc. Some examples of conditions included in the first set of conditions are as follows. In a first example condition, the point belongs to a frame that is not an intra coded frame, or that does not correspond to a random access point. In a second example condition, the azimuth of the current point is different from the azimuth of the previously decoded point(s) of the current frame, or the delta azimuth between the current and previous point is non-zero. In some examples, an approximation of the delta azimuth, or a quantized version of the delta azimuth, may be used to determine the inter-prediction mode of the current node. In a third example condition, the point belongs to a slice that is of a particular type (i.e., P slice or B slice indicating inter-prediction may be applied). In a fourth example condition, a flag is signaled to indicate whether inter prediction is enabled for a particular frame. This flag may be signaled in a syntax element associated with the frame (e.g., slice/slice header, frame header, etc.) or a parameter set to which the frame refers. For instance, an `inter_prediction_enabled_flag` may be signaled in the geometry parameter set (GPS) to specify whether inter prediction is enabled for point cloud frames/slices referring to the GPS. When `inter_prediction_enabled_flag` indicates that inter prediction is not enabled, inter prediction indication for points in the point cloud may not be signaled.

[0094] In some examples, the inter-prediction mode for a point may be determined based on the residual of the azimuth or the phi multiplier associated with the point.

[0095] In some examples, the inter prediction (or one or more inter prediction modes) may be enabled only when the angular mode is enabled (i.e., the inter prediction may be disabled where the angular mode is not enabled).

[0096] The G-PCC coder may chose a reference frame. In some examples, the G-PCC coder may use a previously decoded frame (or in some cases the preceding frame in decoding order) as the reference frame. In other examples, an indication of a frame number (using LSB values or delta frame number values) may be used to specify a reference frame. More generically, two or more frames may be specified as reference frames and a point may be coded with inter-prediction from any of the reference frames (indication of the reference frame associated with the point may be signaled or derived).

[0097] In another example, prediction may be performed from two or more frames (e.g., bi-prediction). As such, a G-PCC coder may predict a point of a current frame based on a reference point in a first reference frame and a reference point in a second reference frame.

[0098] The G-PCC coder may utilize a plurality of inter-prediction modes. When a point is inter-predicted, there may be one or more ways to predict the point from a reference frame. Different mode values may be used to specify each type of prediction (e.g., a G-PCC coder may signal a syntax element that indicates an inter prediction mode used for a current point). As one example inter prediction mode, a point may be predicted from a zero-motion candidate from a reference frame (e.g., a reference point in the reference frame may be a zero-motion candidate). As another example inter prediction mode, a point may be predicted from a global motion candidate from a reference frame. As another example inter prediction mode, a point may be predicted from a candidate point from a reference frame, and other parameters (e.g., motion vector, etc.) may be used to specify the candidate point.

[0099] Some details about example different types of inter-prediction are provided below. The G-PCC coder may utilize a partial inter-prediction mode. In one example, inter-prediction may only apply to predict a subset of the characteristics of a point in a point cloud. One of more of the characteristics for which intra prediction does not apply may be coded using intra prediction or by other techniques. In one example, when a point is indicated to be inter-predicted, the radius of the point may be predicted from a reference point (or a reference frame) when angular mode is enabled, and the azimuth and laser ID (or elevation) may be derived based on intra prediction or other techniques. In another example, if a point is indicated to be inter-predicted, both the radius and the azimuth of

the point are predicted from a reference point (or a reference frame) when angular mode is enabled; the G-PCC coder may derive laser ID (or elevation) based on intra prediction or other techniques.

[0100] The G-PCC coder may perform signaling of the inter-prediction mode. In some examples, the inter-prediction mode associated with a point may be signaled in the bitstream. This may be signaled as a flag, or as a mode value. In one example inter-prediction mode signaling technique, an inter-pred-flag equal to 0 may indicate that the point is predicted using intra-prediction, and inter-pred-flag equal to 1 may indicate that the point is predicted using inter-prediction. In some examples, based on the indication that the point is predicted using inter-prediction, a mode value may be signaled to specify the process to perform the inter-prediction of the point. In another example inter-prediction mode signaling technique, when a point is predicted with inter-prediction, a mode value may be signaled where a mode value of 0 specifies that the zero-motion candidate is used for prediction, and a mode value of 1 specifies that a global motion candidate is used for prediction.

[0101] In this disclosure, inter-prediction mode may be used interchangeably with an indication of whether a point is coded with inter-prediction and with the particular inter-prediction technique used.

[0102] Certain conditions may affect how the G-PCC coder signals the inter-prediction mode. In some cases, inter-prediction mode may be signaled only when a second set of conditions apply. The second set of conditions may be derived using one or more characteristics of the point such as cartesian position, angular coordinates (radius, azimuth, elevation), prediction mode of neighboring nodes, etc. An example condition included in the second set of conditions may include that the point belongs to a frame that is not an intra coded frame, or that does not correspond to a random access point. An example condition included in the second set of conditions may include that the azimuth of the current point is different from the azimuth of the previously decoded point(s), or the delta azimuth between the current and previous point is non-zero. In some examples, an approximation of the delta azimuth, or a quantized version of the delta azimuth, may be used to determine the inter-prediction mode of the current node. An example condition included in the second set of conditions may include that the delta azimuth between the current and the previous point is more than a threshold; the value of the threshold may be fixed or signaled in the bitstream, or derived based on syntax elements signaled in the bitstream.

[0103] More generally, the second set of conditions may include one or more criteria, including but not limited to, a delta azimuth value, a delta radius value, a delta laserID value, a delta elevation value. This delta value of a characteristic may be indicated by a corresponding residual value that is signaled in the bitstream for a point. The second set of conditions may also include delta coordinate values in the Cartesian domain (delta x, delta y, delta z).

[0104] When an inter-prediction mode is not signaled for a particular point, the inter-prediction mode may be inferred to be a fixed value for the point. When inter-prediction mode is not signaled for a point, the value may be assumed to be 0 (or the value that specifies intra prediction).

[0105] Signaling optimizations will now be discussed. In some cases, an inter-prediction mode may not be signaled for all points, and instead the inter-prediction mode may be derived from a syntax element signaled in the bitstream. An inter-prediction mode may be signaled for a prediction tree, and the inter-prediction mode may apply to all the points in the prediction tree. An inter-prediction mode may be signaled in a geometry slice, and the inter-prediction mode may apply to all the points in the slice. An inter-prediction mode may be signaled for a prediction block (specified as a fixed number of points, or a signaled number of points, etc.), and the inter-prediction mode may apply to all the points in the prediction block.

[0106] In some examples, the inter-prediction mode may only be signaled to some nodes that are specified by the tree depth. For example, an inter-prediction mode may only be signaled for the root node of a prediction tree.

[0107] In one example, when inter prediction is applied for some characteristics of the point (e.g., radius), and intra prediction is applied to other characteristics of the points (e.g., azimuth, laserID) only a subset of the intra prediction modes may be allowed (e.g., only Mode 0 of intra prediction may be allowed for azimuth, laser ID). In some examples, if only one intra prediction mode is allowed, then mode value is not signaled and inferred at the decoder. In some examples, if a subset (two or more) of intra prediction modes is allowed, a more optimized coding of the mode syntax element may be applied (e.g., fewer bits may be sufficient to code the index to the subset of the modes).

[0108] For each point and corresponding to each prediction mode, the G-PCC coder may derive a prediction value of one or more characteristics of the current point based on the position of one or more points in the reference frame. An example prediction process follows:

1a. When a zero-motion candidate is chosen, a prediction candidate may be chosen as the point with the closest azimuth and laser ID corresponding to a particular value of the azimuth and laser ID (e.g., reconstructed azimuth value and laserID of the current point); the value of radius of the prediction candidate is chosen as the prediction value of radius of the current node.

1b. When a zero-motion candidate is chosen, a prediction candidate may be chosen as the point with the closest azimuth and laser ID corresponding to a particular value of azimuth and laser ID (e.g., a reconstructed azimuth value and laserID of a previously decoded and reconstructed point); the values of radius, azimuth and laserID of the prediction candidate are chosen as the prediction value of radius, azimuth and laserID of the current node.

2a. When a global motion candidate is chosen, a prediction candidate may be chosen as the point with the closest azimuth and laser ID corresponding to a particular value of azimuth and laser ID after the application of global motion compensation to the prediction candidate; the value of radius of the prediction candidate is chosen as the prediction value of radius of the current node.

2b. When a global motion candidate is chosen, a prediction candidate after the application of global motion compensation to the prediction candidate may be chosen as the point with the closest azimuth and laser ID corresponding to a particular value of azimuth and laser ID (e.g., reconstructed azimuth value and laserID of a previously decoded and reconstructed point); the values of radius, azimuth and laserID of the prediction candidate are chosen as the prediction value of radius, azimuth and laserID of the current node.

3a. When a generic candidate is chosen, a prediction candidate may be chosen as the point with the closest azimuth and laser ID corresponding to a particular value of azimuth and laser ID after the application of motion vector (associated with the current point) to the prediction candidate; the value of radius of the prediction candidate is chosen as the prediction value of radius of the current node.

3b. When a generic candidate is chosen, a prediction candidate after the application of motion vector to the prediction candidate may be chosen as the point with the closest azimuth and laser ID corresponding to a particular value of azimuth and laser ID (e.g., reconstructed azimuth value and laserID of a previously decoded and reconstructed point); the values of radius, azimuth and laserID of the prediction candidate are chosen as the prediction value of radius, azimuth and laserID of the current node.

[0109] In some examples, a motion vector may be used to derive a point that is used to derive a radius, an azimuth and laser ID (or elevation) of the current point, or a prediction of these values. In some examples, the motion vector may be used to derive the prediction of x, y, z coordinates, or a prediction of these values.

[0110] In some examples, more than one point in the reference frame may be used to derive the prediction value. The number of points can be signaled in slice header or in the GPS. In some examples, an index to the two or more points in the reference frame may be signaled to specify the particular point to be used for inter prediction. When a zero-motion candidate is chosen, two prediction candidates are chosen as the two points with the two closest azimuth and laser ID corresponding to a particular value of azimuth and laser ID; the average value of radii of the two prediction candidates is chosen as the prediction value of the radius of the current node. When two or more prediction candidates are chosen as candidates for inter prediction (e.g., based on the closeness of laser ID and/or azimuth), the prediction candidate may be chosen by signaling an index to a list comprising the two or more prediction candidates. In some examples, a weighted average (weights calculated based on the difference in the laser ID values of the current point and predictor and/or a difference between the azimuth of the current point and the predictor) may also be used as a prediction candidate.

[0111] The G-PCC coder may perform preparation of the reference frame. In some examples, the reference frame may be motion compensated before inter prediction. For example, one or more points of a reference frame may be associated with a motion (e.g., global motion), and the prediction for a subsequent frame may be obtained by compensating the position of the points with the associated motion. This motion may be estimated or obtained from external information (such as the GPS information). The “reference frame” used for inter-prediction may thus be a “compensated reference frame.” In other words, the reference frame may be a motion compensated reference frame.

[0112] In some examples, one or more points in the compensated reference frame may be considered as a zero-motion candidate for the current frame. In some examples, the points in the reference frame for which motion compensation is applied may belong to a particular feature in the point cloud (e.g., an object), or it may be labelled (e.g., ground, non-ground) by some estimation algorithms. The global motion parameters (e.g., rotation, translation) may be signaled in the bitstream.

[0113] In some examples, in addition to global motion parameters, additional adjustment values may be signaled for one or more of the following: x, y, z, radius, azimuth, laserID.

For example, when the predictor frame is prepared by applying the motion compensation (e.g., using motion parameters), one or more additional adjustments values of x, y, z may be applied in the Cartesian domain; one or additional adjustments values of radius, azimuth, laserID may be applied in the spherical domain.

[0114] In some examples, the adjustment values may be specified for the entire sequence (representing a “global” adjustment); these values may also be incorporated in the global motion parameters.

[0115] In some examples, the adjustment values may be specified for a group of frames (representing a “group” adjustment); these values may also be incorporated in the global motion parameters.

[0116] In some examples, the adjustment values may be specified for a frame (representing a “frame” adjustment); these values may also be incorporated in the global motion parameters.

[0117] In some examples, the adjustment values may be specified for different regions of the points cloud (representing a “local” adjustment); these values may also be incorporated in the global motion parameters or signaled as modifications to global motion applicable to the frame.

[0118] One or more of global, frame and local adjustments may be signaled; when more than one of these parameters is present, the corresponding adjustments may be applied in succession or together (e.g., a global, frame and local x adjustment x_g , x_f , x_l , respectively may be applied as $x_g + x_f + x_l$ for points belonging to the region). In other examples, all the adjustments at one level (e.g., global) may be applied before applying the adjustments at a different level (e.g., frame)

[0119] The motion vector associated with inter-prediction mode of a point may be signaled in the bitstream. In some cases, multiple points may share the same motion vector and may be signaled in the slice or in a parameter set (e.g., GPS, or parameter set specifically used for motion parameters). In some cases, the motion vector associated with a particular point can be predicted from spatial or temporal inter-predicted neighboring points, and only the difference between the actual and the predicted motion vector can be signaled.

[0120] The G-PCC coder may perform context selection. For instance, the G-PCC coder may determine the context to be used for predictive geometry parameters/syntax elements based on one or more syntax elements signaled in the bitstream, or variables derived

thereof. The use of additional contexts may improve the compression efficiency of the codec.

[0121] In some examples, the G-PCC coder may utilize a first parameter associated with a point to determine the context for other syntax elements associated with the point. For instance, the context of one or more syntax elements associated with phi multiplier (e.g., ptn_phi_mult_abs_gt0_flag, ptn_phi_mult_sign_flag, ptn_phi_mult_abs_gt1_flag, ptn_phi_mult_abs_minus2, ptn_phi_mult_abs_minus9) may be selected based on an inter flag.

[0122] In one example, the inter flag indicates whether a first set of contexts is selected in case inter prediction was used or whether a second set of contexts is selected in case intra prediction was used. This may apply to the coding of the phi multiplier, primary and secondary residuals in case of predictive geometry with angular mode enabled (spherical position coding: azimuth, radius, elevation or laserID).

[0123] If the context for a syntax element is dependent on a first set of conditions, one context may be chosen when the first set of conditions is true and the first parameter associated with the point takes a first value, and a second context may be chosen when the first set of conditions is false, and the first parameter associated with the point takes a second value that is different from the first value.

[0124] Different choices of first parameter are described. Each choice may present a different trade-off between compression efficiency and storage (more contexts may require more storage associated with context states). In one example, the first parameter may be an inter flag associated with the point. In another example, the first parameter may be a prediction mode (e.g., prediction mode used in predictive geometry) associated with the point. In a further example, the first parameter may be a tuple consisting of inter flag and prediction mode associated with a point e.g., (interFlag, predMode). When more than one inter prediction candidate is present (e.g., zero motion candidate, global motion candidate, etc.), let interPredMode represent the various inter prediction candidates (e.g., 0 for intra, 1 for zero motion candidate, 2 for global motion candidate, etc.). The first parameter may be interPredMode.

[0125] In some examples, the G-PCC coder may use different conditions/parameters for different type of frames. For example, for frames that are intra coded (that do not have reference to other frames in the sequence for prediction), a second parameter may be chosen to determine the context as described in the disclosure, and other frames (e.g., inter-coded frames), a third parameter may be chosen to determine the context as

described in this disclosure. The second and third parameters may be different (e.g., for intra frames, the prediction mode may be chosen as the second parameter and for inter frames, the tuple of prediction mode and inter flag may be chosen as the third parameter). The G-PCC coder may extend similar context selection to one or more syntax elements/components (primary residual, secondary residuals, etc.).

[0126] FIG. 7 is a conceptual diagram illustrating an example inter prediction process for predicting points of a point cloud, in accordance with one or more aspects of this disclosure. As shown in FIG. 7, current frame 750 may include a plurality of points 752A-752L (collectively, “points 752”) and reference frame 754 may include a plurality of points 756A-756L (collectively, “points 756”). Reference frame 754 may be a frame that is encoded and/or reconstructed prior to current frame 750 being decoded and/or reconstructed (e.g., reference frame 754 may precede current frame 750 in a coding order). A G-PCC coder may utilize intra prediction to predict one or more of points 752 of current frame 750 based on one or more of points 756 of reference frame 754. For instance, a G-PCC decoder (or a reconstruction loop of a G-PCC encoder) may predict one or more parameters (e.g., (r, ϕ, i)) of current point 752A of points 752 based on one or more of points 756.

[0127] To perform inter prediction to predict a current point in a current frame, the G-PCC coder may determine a reference point in a reference frame that is different than the current frame and predict one or more parameters of the current point based on the reference point. For instance, to predict current point 752A, the G-PCC coder may determine reference point 756A and predict one or more parameters of current point 752A based on one or more parameters of reference points 756A. The determined reference point may be referred to as an identified reference point.

[0128] The G-PCC coder may determine the reference point using any suitable technique. As one example, the G-PCC coder may determine, in the current frame, a pivot point that precedes the current point in a coding order; and determine, based on one or more parameters of the pivot point, the reference point. For instance, where the coding order is counter-clockwise, the G-PCC coder may determine that point 752B is the previous point of current point 752A (e.g., the point that immediately precedes the current point in the coding order) (i.e., determine that point 752B is the pivot point), and determine a reference point based on one or more parameters of pivot point 752B.

[0129] To determine the reference point based on the one or more parameters of the pivot point, the G-PCC coder may determine, in the reference frame, a reference pivot point

based on an azimuth of the pivot point; and determine the reference point based on the reference pivot point. For instance, the G-PCC coder may determine, in reference frame 754, a point having a same azimuth (or same azimuth and same laser ID) as pivot point 752B. In the example of FIG. 7, the G-PCC coder may determine that point 756B is the reference pivot point as point 756B has a same azimuth as pivot point 752B. While the reference pivot point in the example of FIG. 7 corresponds to an actual point (e.g., an actual point in frame 754), the techniques of this disclosure are not necessarily so limited. For instance, in some examples, the reference pivot point may be a virtual point that does not correspond to a reconstructed point in reference frame 754.

[0130] In some examples, the G-PCC coder may determine the reference pivot point based on an actual (e.g., unscaled) azimuth of the pivot point. In other examples, the G-PCC coder may determine the reference pivot point based on a scaled azimuth of the pivot point. For instance, the G-PCC coder may determine the scaled azimuth by scaling the azimuth of the pivot point by a constant value.

[0131] To determine the reference point based on the reference pivot point, the G-PCC coder may identify, as the reference point in the reference frame, a point having an azimuth that greater than an azimuth of the reference pivot point. For instance, the G-PCC coder may determine which of points 756 have azimuth values greater than an azimuth value of the reference pivot point, and select (from the set of points 756 having azimuth values greater than the azimuth value of the reference pivot point) the point having the smallest azimuth value. In this example of FIG. 7, point 756A may be the point in reference frame 754 that has the smallest azimuth that is larger than the azimuth of reference pivot point 752B. As such, the G-PCC coder may identify point 756A as the reference point for performing intra prediction of current point 752A.

[0132] In some examples, the G-PCC coder may determine the reference point based on an actual (e.g., unscaled) azimuth of the reference pivot point. In other examples, the G-PCC coder may determine the reference point based on a scaled azimuth of the reference pivot point. For instance, the G-PCC coder may determine the scaled azimuth of the reference pivot point by scaling the azimuth of the pivot point by a constant value. As such, in some examples, the G-PCC coder may determine the reference point by identifying, as the reference point, a point having a smallest scaled azimuth that is larger than the scaled azimuth of the reference pivot point (e.g., point 756A). In some examples, the G-PCC coder may utilize the point with the second smallest azimuth greater than the scaled azimuth. For instance, the G-PCC coder may determine the reference point by

identifying, as the reference point, a point having a second smallest scaled azimuth that is larger than the scaled azimuth of the reference pivot point (e.g., point 756L).

[0133] The G-PCC coder may predict parameters of current point 752A based on parameters of reference point 756A. For instance, the G-PCC coder may signal a residual data that represents differences between parameters of current point 752A and reference point 756A. The G-PCC decoder may add the residual data to the parameters of reference point 756A to reconstruct the parameters of current point 752A.

[0134] While discussed above as using a single reference point in a single reference frame, the techniques of this disclosure are not so limited. As one example, multiple reference points in a single reference frame may collectively be used to predict the current point. For instance, the G-PCC coder may determine, in a reference frame and based on a reference pivot point, a plurality of reference points. The G-PCC coder may predict one or more parameters of a current point in a current frame based on the plurality of reference points. As another example, reference points from multiple reference frames may be used to predict the current point.

[0135] As discussed above, a G-PCC coder may perform azimuth prediction. Let (r , ϕ , laserID) be the three coordinates of the pivot point (referred to as radius, azimuth and laser ID) in the spherical coordinate system. Techniques disclosed herein may also apply to other coordinate systems.

[0136] In some examples, the G-PCC coder may code points in the current point cloud frame in an ordered fashion as follows:

[0137] First, for a current point in the current frame, the G-PCC coder may choose a pivot point in the current frame that is preceding the first point in decoding order. In some examples, the pivot point is the preceding point in the current frame in the decoding order. In some examples, the pivot point is the second preceding point in the current frame in the decoding order. More generally, more than one preceding point may be chosen as pivot points of the current point. In some examples, the pivot point may be a virtual point derived based on a previously decoded point in the current frame and azimuth displacement that are a multiple of azimuth quantization scale value (pre-determined or derived from signaled syntax elements).

[0138] Second, the G-PCC coder may choose a point in reference frame, reference pivot point, that is associated with the pivot point. The reference pivot point may be chosen as a point in the reference frame that has the same azimuth and laser ID as the pivot point. In some examples, points with other laser ID values may also be candidates for reference

pivot point (e.g., the reference pivot point may be chosen as the point in the reference frame that has the same azimuth as the pivot point and laser ID in the range of $[\text{LaserID} - M, \text{LaserID} + M]$, where LaserID is the laser ID of the pivot point and M is a fixed value (e.g., 1) or chosen based on distance of the pivot point from the origin, or derived as function of LaserID (e.g., for smaller values of LaserID, M may be smaller and for larger values of LaserID, M may be larger)). In some examples, a distance metric may be defined using azimuth and laser ID, and the reference pivot point is chosen as the point that has smallest distance from the azimuth and laser ID of the pivot point using the distance metric. A normalized azimuth value may be obtained by scaling the azimuth by a first constant value; a normalized laser ID may be obtained by scaling the laserID by a second constant value; the distance metric may be obtained by computing a norm (e.g., L2-norm, L1-norm) for the normalized laser ID and azimuth value of the pivot point and the reference points. More generally, the reference pivot point may be chosen as a point in the reference frame that is in the neighborhood of points with the same azimuth and laser ID as the pivot point. In some examples, the reference pivot point may be a virtual point in the reference frame that is derived from the pivot point that has the same azimuth and laser ID as the pivot point.

[0139] Third, the G-PCC coder may choose a reference point in the reference frame that is associated with the reference pivot point. The reference point may be chosen as the point in the reference frame that has the smallest azimuth that is larger than that of reference pivot point and the same laser ID as the reference pivot point. The reference point may be chosen as the point in the reference frame that has the second smallest azimuth that is larger than that of reference pivot point and the same laser ID as the reference pivot point. In some examples, when the reference point is not available, inter prediction may be disabled for the current point. In some examples, the reference point may be chosen as the reference pivot point.

[0140] Fourth, the G-PCC coder may compute a first residual between the reference point and the reference pivot point.

[0141] Fifth, the G-PCC coder may use the first residual to derive a first prediction for the current value. The prediction may be derived by adding the components of the first residual to the respective components of the pivot point (e.g., the radius prediction may be obtained by adding the radius component of the first residual to the radius component of the pivot point (similarly for azimuth)). In some examples, the first prediction may be set equal to the reference point.

[0142] Sixth, the G-PCC coder may code a second residual between the first prediction and the position of the current point.

[0143] Seventh, composition of residual: one or more residuals disclosed in this disclosure may include one or more of the following: The residual may comprise a radius residual between the reference pivot point and the reference point. The residual may comprise an azimuth residual between the reference pivot point and the reference point.

[0144] Eighth, the G-PCC coder may derive a current point based on the first prediction and second residual. In some examples, the G-PCC coder may derive the current point from the second residual (e.g., and not based on the first prediction).

[0145] The G-PCC coder may apply one or more techniques described above to a quantized azimuth value; the scale value used for the quantization may be derived from signaled values or pre-determined. The quantized azimuth value and laserID may be used to search an inter predicted point in the reference. For example, the azimuth of a previously decoded and reconstructed point may be quantized and an inter predicted point with a quantized azimuth and laserID that is closest to the quantized azimuth and laserID of the previous point may be selected as predictor for the azimuth, radius and laserID of the current point, either as quantized/dequantized predictor or unquantized.

[0146] The G-PCC coder may apply one or more techniques described above to a quantized laserID value; the scale value used for the quantization may be derived from signaled values or pre-determined. The azimuth and quantized laserID value may be used to search an inter predicted point in the reference. For example, the laserID of a previously decoded and reconstructed point may be quantized and an inter predicted point with an azimuth and quantized laserID that is closest to the azimuth and quantized laserID of the previous point may be selected as predictor for the azimuth, radius and laserID of the current point, either as quantized/dequantized predictor or unquantized.

[0147] The G-PCC coder may apply one or more techniques described above to a quantized azimuth and quantized laserID value; the scale values used for the quantization may be derived from signaled values or pre-determined. The quantized azimuth and quantized laserID value may be used to search an inter predicted point in the reference. For example, the azimuth and laserID of a previously decoded and reconstructed point may be quantized and an inter predicted point with a quantized azimuth and quantized laserID that is closest to the quantized azimuth and quantized laserID of the previous point may be selected as predictor for the azimuth, radius and laserID of the current point, either as quantized/dequantized predictor or unquantized.

[0148] In some examples, the reference frame may refer to a set of (radius, azimuth, laserID) tuples that are derived from the reference frame. For example, for each point in the reference frame, the G-PCC coder may add a radius, azimuth, laser ID to the set if there is no other point in the set that has the same azimuth the laser ID. In some cases, a quantized value of azimuth may be added. In some cases, a (r, phi, laserID) may be added even if there is another tuple in the set with same phi and laserID, e.g., (r1, phi, laserID), if the value of r is less than the value or r1 (in this case, the tuple (r1, phi, laserID) that was present may be replaced by the new (r, phi, laserID)). In some cases, the points in the reference frame may be in the x,y,z domain; the points may be stored for reference as is, or by converting it to the spherical domain. In some cases, a motion compensated position may be added to the reference frame. The compensation may be based on a motion vector (e.g., global motion vector with rotation and/or translation) signaled that is associated with the current frame and reference frame.

[0149] The G-PCC coder may perform context selection for inter-prediction modes. As discussed above, the G-PCC coder may signal inter prediction flag and inter prediction mode. The context to be used by the G-PCC coder for coding the inter prediction flag or inter prediction mode may be chosen as follows: The G-PCC coder may set b1, b2, b3, ... bN to be the inter prediction flag values for the N previously decoded points (b1 is the immediately preceding node, b2, immediately preceding b1, etc. The G-PCC coder may chose the context index as the number generated by using the b1, b2, ... bN as follows: $ctxIdx = b1 + b2 \ll 1 + b3 \ll 2 + \dots bN \ll (N-1)$. In another example, the G-PCC coder may chose the context index as $b1 + b2 + \dots + bN$. In some examples, the value of b1, b2, .. bN, may also be used to choose the context of other syntax elements such as inter prediction mode, motion vector, etc. In some examples N could be fixed as 5.

[0150] The G-PCC coder may perform improved radius residual coding. The radius, azimuth and laserID associated with each point may be coded in predictive geometry coding. The residual of the radius component of each point may be coded with an "equal_to_zero" flag, a sign bit, a number of bits for the remaining residual and/or the remaining residual. The number of bits may be an indicator of the relative magnitude of the radius residual. When LIDAR captured points are close to each other in the point cloud (e.g., points of an object or building captured by one of the lasers), the radius may not change by much across the points. Since the object (most likely solid) is continuous, successive LIDAR scan points may be returned to the sensor where the azimuth difference between successive points is close to the sampling frequency of the azimuth. The

difference in the azimuth, or azimuth residual, for adjacent points (of the same laser) on an object is close to zero.

[0151] The context of number of bits of radius residual is chosen based on the azimuth residual of the current point. For example, based on the magnitude of the azimuth residual the context may be switched as follows where N is a fixed value (signaled or predetermined). Table 1, below, provides one example of context index for number of bits used in number of bits for radius residual N.

Table 1: Context index for number of bits used in number of bits for radius residual

Abs(azimResi)	N=3	N=4
0	1	1
1	1	1
2	2	2
3	3	3
4	0	4
>=5	0	0

[0152] Note that other processes to choose the context index based on the radius residual may also be used.

[0153] The following specific examples show implementations of several techniques of this disclosure.

Example A

[0154] For each point in the point cloud frames that are not all-intra coded (random access point), a flag may be used to specify whether a point is coded with inter-prediction. When inter-prediction is applied, a “zero motion vector” candidate is chosen to predict the radius of predicted point from a point in a reference frame that has the same laser ID value as the current point and that has the closest quantized azimuth value as the current point. Changes to the syntax structure is provided below (<ADD>...</ADD> tags indicate additions and <DELETE> ... </DELETE> tags indicate deletions):

geometry_predtree_node(nodeIdx) {	Descriptor
if(geom_scaling_enabled_flag && !(nodeIdx % PtnQpInterval)) {	
ptn_qp_offset_abs_gt0_flag	ae(v)
if(ptn_qp_offset_abs_gt0_flag) {	
ptn_qp_offset_sign_flag	ae(v)
ptn_qp_offset_abs_minus1	ae(v)
}	

}	
if(unique_geometry_points_flag) {	
ptn_point_cnt_gt1_flag	ae(v)
if(ptn_point_cnt_gt1_flag)	
ptn_point_cnt_minus2	ae(v)
}	
ptn_child_cnt[nodeIdX]	ae(v)
ptn_pred_mode[nodeIdX]	ae(v)
if(geometry_angular_enabled_flag) {	
ptn_phi_mult_abs_gt0_flag	ae(v)
if(ptn_phi_mult_abs_gt0_flag) {	
ptn_phi_mult_sign_flag	ae(v)
ptn_phi_mult_abs_gt1_flag	ae(v)
if(ptn_phi_mult_abs_gt1_flag)	
ptn_phi_mult_abs_minus2	ae(v)
if(ptn_phi_mult_abs_minus2 == 7)	
ptn_phi_mult_abs_minus9	ae(v)
}	
}	
<ADD>if(PtnPhiMult[nodeIdx] && InterEnableFlag)	
ptn_inter_flag[nodeIdX]	u(1) </ADD>
numComp = geometry_angular_enabled_flag && !number_lasers_minus1 ? 2 : 3	
for(k = 0; k < numComp; k++) {	
ptn_residual_abs_gt0_flag[k]	ae(v)
if(ptn_residual_abs_gt0_flag[k]) {	
if(k ptn_pred_mode[nodeIdX] <ADD> ptn_inter_flag[nodeIdX] </ADD>)	
ptn_residual_sign_flag[k]	ae(v)
ptn_residual_abs_log2[k]	ae(v)

ptn_residual_abs_remaining[k]	ae(v)
}	
}	
if(geometry_angular_enabled_flag)	
for(k = 0; k < 3; k++) {	
ptn_sec_residual_abs_gt0_flag[k]	ae(v)
if(ptn_sec_residual_abs_gt0_flag[k]) {	
ptn_sec_residual_sign_flag[k]	ae(v)
ptn_sec_residual_abs_gt1_flag[k]	ae(v)
if(ptn_sec_residual_abs_gt1_flag[k])	
ptn_sec_residual_abs_minus2[k]	ae(v)
}	
}	
for(i = 0; i < ptn_child_cnt[nodeId] ; i++)	
geometry_predtree_node(++PtnNodeIdx)	
}	

<ADD>

ptn_inter_flag[nodeId] equal to 0 specifies that the radius residual of the current node is coded with intra-prediction. **ptn_inter_flag[nodeId]** equal to 1 specifies that the radius residual of the current node is coded with inter-prediction. When not present, the value of **ptn_inter_flag[nodeId]** is inferred to be equal to 0.

</ADD>

[0155] The value of **InterEnableFlag** is set to 0 when the current point cloud frame is a random access point or an all-intra frame, and set to 1 otherwise.

[0156] In some examples, when **ptn_inter_flag** is equal to 1, the syntax element **ptn_pred_mode[]** is not signaled and is inferred to be a default value (e.g., equal to 0 corresponding to predictive geometry mode 0).

[0157] Changes to the position prediction process are as follows:

Position prediction process

The inputs to this process are:

- the variable predMode indicating the prediction mode of current node,
- the arrays aPos0, aPos1, and aPos2, with values aPosX[k], k = 0 .. 2, and X = 0 .. 2. Each array contains the position associated with the X-th generation ancestor node in the depth-first tree traversal order,
- the variable curDepth, indicating the distance in nodes between the current node and the root node of the current predictive tree,

<ADD>

- A variable interFlag indicating whether the radius of the current node is coded with inter-prediction.

</ADD>

The output from this process is the array predPos, with values predPos[k], k = 0 .. 2, indicating the predicted point position associated with the nodeIdx-th tree node.

When predMode is equal to 0, the predicted point position is calculated as follows,

for (k = 0; k < 3; k++)

predPos[k] = 0

if (geometry_angular_enabled_flag) {

predPos[0] = ptn_radius_min_value

if (curDepth > 0) {

predPos[1] = aPos0[1]

predPos[2] = aPos0[2]

}

}

When predMode is equal to 1, the predicted point position is the position associated with the first ancestor.

for (k = 0; k < 3; k++)

predPos[k] = aPos0[k]

When predMode is equal to 2, the predicted point position is a linear combination of the positions associated with the first two ancestors

for (k = 0; k < 3; k++)

$$\text{predPos}[k] = \text{aPos0}[k] + \text{aPos0}[k] - \text{aPos1}[k]$$

Otherwise, predMode is equal to 3, the predicted point position is a linear combination of the positions associated with all three ancestors

for (k = 0; k < 3; k++)

$$\text{predPos}[k] = \text{aPos0}[k] + \text{aPos1}[k] - \text{aPos2}[k]$$

<ADD>

When interFlag is equal to 1, predPos[0] is derived as follows. Let refFramePos[k] for k = 0..2 be the position of the point in the reference frame such that refFramePos[2] = predPos [2] and the absolute difference between the value of refFramePos[1] and predPos [1] is the lowest among the points in the reference frame. predPos[0] is set equal to refFramePos[0].

</ADD>

In one example, a quantized value of refFramePos[1] and predPos [1] is used to measure the absolute difference in values, where the quantization scale may be a fixed number (e.g., geom_angular_azimuth_step_minus1 + 1)

Example B

[0158] In example B, the prediction mode may not be signaled when inter prediction is used. For instance, when inter prediction is used, prediction mode is inferred to be equal to 0. Changes to the syntax structure is provided below (<ADD>...</ADD> tags indicate additions and <DELETE> ... </DELETE> tags indicate deletions):

}	
ptn_child_cnt[nodeId]	ae(v)
<DELETE> ptn_pred_mode[nodeId]	ae(v) </DELETE> >
if(geometry_angular_enabled_flag) {	
ptn_phi_mult_abs_gt0_flag	ae(v)
if(ptn_phi_mult_abs_gt0_flag) {	
ptn_phi_mult_sign_flag	ae(v)
ptn_phi_mult_abs_gt1_flag	ae(v)

if(ptn_phi_mult_abs_gt1_flag)	
ptn_phi_mult_abs_minus2	ae(v)
if(ptn_phi_mult_abs_minus2 == 7)	
ptn_phi_mult_abs_minus9	ae(v)
}	
}	
<ADD>if(PtnPhiMult[nodeIdx] && InterEnableFlag)	
ptn_inter_flag[nodeIdx]	u(1)
if(!ptn_inter_flag[nodeIdx])	
ptn_pred_mode[nodeIdx]	ae(v) </ADD>
numComp = geometry_angular_enabled_flag && !number_lasers_minus1 ? 2 : 3	
for(k = 0; k < numComp; k++) {	

[0159] When ptn_pred_mode[nodeIdx] is not signaled, ptn_pred_mode[nodeIdx] may be inferred to be equal to 0. In some examples, the following inference may be added: When ptn_inter_flag[nodeIdx] is equal to 1, ptn_pred_mode[nodeIdx] is inferred to be equal to 0.

Example C

[0160] In example C, the order of signaling of phi and prediction mode may not be modified. The signaling of prediction mode may be conditioned on the signaling of inter flag. The signaling of ptn_phi_mult_abs_gt0_flag may also be conditioned on the inter flag. Changes to the syntax structure are provided below (<ADD>...</ADD> tags indicate additions and <DELETE> ... </DELETE> tags indicate deletions):

}	
ptn_child_cnt[nodeIdx]	ae(v)
<ADD>if(InterEnableFlag)	
ptn_inter_flag[nodeIdx]	u(1)

if(!ptn_inter_flag[nodeIdX])	
ptn_pred_mode[nodeIdX]	ae(v) </ADD>
if(geometry_angular_enabled_flag) {	
<ADD>if(!ptn_inter_flag[nodeIdX]) </ADD>	
ptn_phi_mult_abs_gt0_flag	ae(v)
if(ptn_phi_mult_abs_gt0_flag) {	
ptn_phi_mult_sign_flag	ae(v)
ptn_phi_mult_abs_gt1_flag	ae(v)
if(ptn_phi_mult_abs_gt1_flag)	
ptn_phi_mult_abs_minus2	ae(v)
if(ptn_phi_mult_abs_minus2 == 7)	
ptn_phi_mult_abs_minus9	ae(v)
}	
}	
numComp = geometry_angular_enabled_flag && !number_lasers_minus1 ? 2 : 3	
for(k = 0; k < numComp; k++) {	

[0161] When ptn_inter_flag[nodeIdX] is not signaled, ptn_inter_flag[nodeIdX] is inferred to be equal to 0.

[0162] When ptn_phi_mult_abs_gt0_flag[nodeIdX] is not signaled, ptn_phi_mult_abs_gt0_flag[nodeIdX] is inferred to be equal to 0. In another example, the following inference is added: When ptn_inter_flag[nodeIdX] is equal to 1, ptn_phi_mult_abs_gt0_flag[nodeIdX] is inferred to be equal to 1.

Example D

[0163] In example D, the order of pred mode and phi syntax elements may not be changed. For instance, inter flag may be signaled after pred mode but before the phi syntax elements.

}	
---	--

ptn_child_cnt[nodeIdIdx]	ae(v)
ptn_pred_mode[nodeIdIdx]	ae(v)
<ADD>if(InterEnableFlag && !ptn_pred_mode[nodeIdIdx])	
ptn_inter_flag[nodeIdIdx]	u(1) </ADD>
if(geometry_angular_enabled_flag) {	
<ADD> if(!ptn_inter_flag[nodeIdIdx]) </ADD>	
ptn_phi_mult_abs_gt0_flag	ae(v)
if(ptn_phi_mult_abs_gt0_flag) {	
ptn_phi_mult_sign_flag	ae(v)
ptn_phi_mult_abs_gt1_flag	ae(v)
if(ptn_phi_mult_abs_gt1_flag)	
ptn_phi_mult_abs_minus2	ae(v)
if(ptn_phi_mult_abs_minus2 == 7)	
ptn_phi_mult_abs_minus9	ae(v)
}	
}	
numComp = geometry_angular_enabled_flag && !number_lasers_minus1 ? 2 : 3	
for(k = 0; k < numComp; k++) {	

[0164] When ptn_pred_mode[nodeIdIdx] is greater than 0 and InterEnableFlag is equal to 1 (ptn_inter_flag[] is not signaled), ptn_inter_flag[nodeIdIdx] is inferred to be 0.

Example E

[0165] Example E is similar to Example A, with the modification that the signaling of inter flag is also dependent on the value of pred_mode. The entirety of the following syntax is added.

if((PtnPhiMult[nodeIdx] !ptn_pred_mode[nodeIdIdx]) && InterEnableFlag)	
--	--

ptn_inter_flag[nodeIdX]	u(1)
----------------------------------	-------------

[0166] When InterEnableFlag is equal 1, and PtnPhiMult[nodeIdX] is equal to 0 or ptn_pred_mode[nodeIdX] is not equal to 0, ptn_inter_flag[nodeIdX] is inferred to be equal to 0.

Example F

[0167] In this example, the prediction for radius, predPos[0], is obtained by not just comparing the azimuth of points that have the same laser ID as the current point, but also points of adjacent laser IDs.

[0168] When interFlag is equal to 1, the G-PCC coder may derive predPos[0] as follows. The G-PCC coder may set refFramePos[k] for $k = 0..2$ to be the position of the point in the reference frame such that refFramePos[2] takes one of values in $\{ \text{predPos}[2] - 1, \text{predPos}[2], \text{predPos}[2] + 1 \}$ and the absolute difference between the value of refFramePos[1] and predPos[1] is the lowest among the points in the reference frame. The G-PCC coder may set predPos[0] equal to refFramePos[0].

[0169] In one case, points from more laser IDs in the reference frame may be considered to obtain the reference point.

[0170] In another case, instead of comparing the azimuth (i.e., refFramePos[1] and predPos[1]), a weighted cost function may be specified as follows:

$$J(\text{refPos}) = w1 * f1(\text{refPos}[1], \text{predPos}[1]) + w2 * f2(\text{refPos}[2], \text{predPos}[2])$$

[0171] The point in the reference frame with position refPos[k] for $k = 0..2$, that is chosen for predicting the radius is the point that has the smallest value of J(refPos).

[0172] In one case, the value of f1(x, y) may be a square of a difference between x and y, or more generally any function that indicates a distance between the current point and reference point in the azimuth (or second) dimension.

[0173] In one case, the value of f2(x, y) may be a square of difference/absolute difference between x and y, or more generally any function that indicates a distance between the current point and reference point in the laser ID (or third) dimension.

[0174] In another case, more than one reference candidate may be chosen for a point and an inter pred candidate index may be signaled to specify which reference candidate is chosen for the point.

[0175] For example, for a current point, the three points in the reference frame may be chosen as follows:

- Point with laser IDs equal to $\text{predPos}[2] - 1$ that has an azimuth value that is closest to $\text{predPos}[1]$
- Point with laser IDs equal to $\text{predPos}[2]$ that has an azimuth value that is closest to $\text{predPos}[1]$
- Point with laser IDs equal to $\text{predPos}[2] + 1$ that has an azimuth value that is closest to $\text{predPos}[1]$

[0176] An index to this set/list may be signaled to specify the point that is chosen for prediction.

[0177] More generally, for each current point that is coded with inter prediction for radius:

- A reference frame is chosen based on a pre-determined decision (e.g., previous decoded frame) or based on one or more signaled indications (e.g., frame index/counter value).
- A set RefCandSuperSet is chosen among points in the reference frame based on a first set of conditions (e.g., one or more processes disclosed herein based on laser ID range and/or azimuth range)
- A subset RefCandSubset is chosen from RefCandSuperSet based on a second set of conditions (e.g., the closest azimuth from the current point for a particular laser ID, etc.)
- If RefCandSubset has more than one entry, arrange the entries in a list RefCandList based on a third set of conditions (e.g., arrange the entries based on increasing laser ID difference from current point); otherwise (only one candidate in RefCandSubset), the one candidate in RefCandSubset is used for prediction.
 - Signal an index to the entry in RefCandList that is used for prediction.
 - Use the prediction entry for predicting the radius of the current point.

[0178] In some examples, when no point in the reference frame exists in RefCandSubset , inter prediction may be disabled, and the inter flag may or may not be signaled in this case.

Example G

[0179] In accordance with one or more aspects of this disclosure, a G-PCC coder may signal, in a bitstream, parameters for deriving the azimuth quantization/scale. In some examples, the G-PCC coder may inter-predict all the three components using a previously decoded point in a current point cloud frame.

[0180] The azimuth scale value (used in deriving the inter prediction candidate) may be signaled in the GPS as follows:

geometry_parameter_set() {	Descriptor
gps_geom_parameter_set_id	u(4)
...	
}	
gps_extension_flag	u(1)
if(gps_extension_flag) {	
inter_prediction_enabled_flag	u(1)
if(inter_prediction_enabled_flag && pred_geom_enabled_flag)	
inter_azim_scale_log2	ue(v)
}	
byte_alignment()	
}	

[0181] In some examples, inter_azim_scale_log2 may be signaled irrespective of inter_prediction_enabled_flag.

[0182] In another examples, the angular mode may also be used to condition signaling of inter_azim_scale_log2.

if (inter_prediction_enabled_flag && pred_geom_enabled_flag && geom_angular_enabled_flag)	
inter_azim_scale_log2	ue(v)

[0183] The following is example semantics for the aforementioned syntax elements.

inter_prediction_enabled_flag equal to 1 specifies that inter prediction may be used for the point cloud frames that refer to the GPS. inter_prediction_enabled_flag equal to 0 specifies that inter prediction is not used for the point cloud frames that refer to the GPS.

inter_azim_scale_log2 is used to scale the azimuth points that may be used in the derivation of inter prediction reference. The value of inter_azim_scale_log2 shall be in the range of 0 to numAzimBits, inclusive.

(The variable numAzimBits may specify a maximum number of bits used to represent the azimuth.)

[0184] In some examples, the inter_azim_scale_log2 syntax element may be an example of a syntax element that specifies a log base 2 value of a scaling factor. A G-PCC coder may determine scaled azimuths (e.g., a scaled azimuth of a pivot point) based on the scaling factor.

[0185] When a frame is coded, a reconstructed spherical position of one or more points in the frame may be stored in a reference frame that may be used as inter prediction reference by a future frame. The reconstructed spherical position may be stored in a hash table as follows:

- The G-PCC coder may obtain a spherical coordinate representation for a reconstructed point.
- The G-PCC coder may obtain a derived azimuth and laser index from the reconstructed point, and use the tuple of derived azimuth and laser index as the index to the hash table.

E.g., the derived a quantized azimuth, denoted as `quantized(val)` may be derived by the G-PCC coder as follows:

```
int offset = azimuthScaleLog2 ? (1 << (azimuthScaleLog2 - 1)) : 0;
quantized(val) = val >= 0 ? (val + offset) >> azimuthScaleLog2
: -((-val + offset) >> azimuthScaleLog2);
```

[0186] Note that in some cases, the `azimuthScaleLog2` may not be signaled and inferred to one of the following:

- The value 0 (i.e., no quantization).
 - The smallest power of 2 that is smaller than or equal to azimuth speed signaled in the bitstream.
 - The smallest power of 2 that is smaller than or equal to (azimuth speed/2) signaled in the bitstream.
- Different sequences may have different azimuth scale values.
- Each entry in the hash table may be derived from one or more reconstructed points.
 - In one example, the hash table entry for a frame may be the first point in decoding order in the reconstructed frame that has a particular value of derived azimuth and laser index.
 - In one example, the hash table entry for a frame may be the point in the reconstructed frame that has the smallest radius among the points that have a particular value of derived azimuth and laser index.
 - In yet another example, the hash table entry for a frame may be the point value that corresponds with the average, mean, median, weighted average, geometric mean, etc. computed using the points that have a particular value of derived azimuth and laser index.

[0187] In some examples, a reconstructed point position in spherical domain (r, phi, laserID) may be entered in the azimuth table as follows: Table index = (quantized(Phi), laserID), Table entry = (r, phi, laserID).

[0188] When coding a current point (in a current frame), the following steps may be performed:

- The G-PCC coder may set prevNode = (r1, phi1, laserID1) to be the previous node in the current frame in the decoding order. In some examples, the prevNode may be chosen as the parent node of the current node.
- In the table associated with the reference frame, the G-PCC coder may check if there exists index (quantized(phi), laserID1) where quantized(phi) is greater than quantized(phi1). If such an index does not exist, the G-PCC coder may not apply inter prediction for the current node. If at least one such index exists, the G-PCC coder may choose the index with the smallest quantized(phi) (which is greater than quantized(phi1)).
- The G-PCC coder may use the entry associated with the above chosen index (r_inter, phi_inter, laserID_inter) as the inter predictor for the current point.

[0189] In some examples, the following code may be used to implement the above technique with a combination of array of hash tables (instead of a 2D hash table), where computePhiQuantized() corresponds to the quantize() function described earlier. In the case below, refPointVals is the vector of hash tables that stores the reference points, refPointVals[currLaserId] is the hash table that corresponds to index laserID = currLaserId, and all entries in refPointVals[currLaserId] are stored with increasing values of quantized(phi). Because this table is stored with increasing quantized(phi), the upper_bound function is used to search the first quantized(phi) laser is greater than the quantized current azimuth quantize(currAzim). In the example below, idx equal refPointVals[currLaserId].end(), there is no such table index and hence inter prediction is not used. When present, idx->second specifies the inter prediction candidate (r_inter, phi_inter, laserID_inter).

[0190] In some examples, the hash table may not store the points with increasing value of quantized(phi), and instead use a hash function to parse and recover table entries. In another example, a generic table/data structure may be used to store the points from the reference frame. As noted above, the following code may be used to implement the above technique.

```
auto quantizedPhi = computePhiQuantized(currAzim);
```

```

if (refPointVals[currLaserId].size()) {
    auto idx = refPointVals[currLaserId].upper_bound(quantizedPhi);
    if (idx == refPointVals[currLaserId].end())
        return std::pair<bool, point_t>(false, 0);
    else
        return std::pair<bool, point_t>(true, idx->second);
}
return std::pair<bool, point_t>(false, 0);

```

- If inter prediction is used, the G-PCC coder may signal a flag indicating inter prediction is used.
- Use the inter prediction candidate derived above to code the residual (at the decoder the prediction candidate is used to add the residual and derive the reconstructed spherical coordinate position of the current point).

[0191] Predictive geometry coding uses a prediction tree structure to predict the positions of the points. When angular coding is enabled, the x, y, z coordinates may be transformed to radius, azimuth and laserID and residuals may be signaled in these three coordinates as well as in the x, y, z dimensions. The intra prediction used for radius, azimuth and laserID may be one of four modes and the predictors are the nodes that are classified as parent, grand-parent and great-grandparent in the prediction tree with respect to the current node. The predictive geometry coding, as currently designed in G-PCC Ed.1, is an intra coding tool as it uses points in the same frame for prediction. Additionally, using points from previously decoded frames may provide a better prediction and thus better compression performance.

[0192] Inter prediction has been proposed to predict the radius of a point from a reference frame. For each point in the prediction tree, a G-PCC coder may determine whether the point is inter predicted or intra predicted (e.g., indicated by a flag). When intra predicted, the intra prediction modes of predictive geometry coding may be used. When inter-prediction is used, the azimuth and laserID are still predicted with intra prediction, while the radius may be predicted from the point in the reference frame that has the same laserID as the current point and an azimuth that is closest to the current azimuth. In some examples, a G-PCC coder may perform inter prediction of the azimuth and laserID (e.g., in addition to radius prediction). When inter-coding is applied, the G-PCC coder may predict the radius, azimuth and laserID of the current point based on a point that is near

the azimuth position of a previously decoded point in the reference frame. In addition, separate sets of contexts may be used for inter and intra prediction.

[0193] As discussed above, FIG. 7 illustrates an example of inter-prediction of current point (curPoint) from a point (interPredPt) in the reference frame. The extension of inter prediction to azimuth, radius, and laserID may include of the following steps:

- For a given point, the G-PCC coder may choose the previous decoded point (prevDecP0).
- The G-PCC coder may choose position in reference frame (refFrameP0) that has same scaled azimuth and laserID as prevDecP0.
- In reference point cloud frame, the G-PCC coder may find the first point (interPredPt) that has azimuth greater than that of refFrameP0. The interPredPt may also be referred to as the “Next” inter predictor.

[0194] FIG. 8 is a flowchart illustrating an example decoding flow associated with the “inter flag” that is signaled for every point. The inter flag signaled for a point indicates whether inter prediction is applied for the point. The flowcharts of this disclosure are provided as examples. Other examples may include more, fewer, or different steps, or steps may be performed in different orders.

[0195] In the example of FIG. 8, G-PCC decoder 300 may determine whether an inter flag of a next point to be decoded (i.e., a current point of a current frame of point cloud data) indicates that the current point is inter predicted (800). If the inter flag of the current point does not indicate that the current point is inter predicted (“NO” branch of 800), G-PCC decoder 300 may identify an intra prediction candidate (812). For instance, G-PCC decoder 300 may determine an intra prediction strategy (e.g., no prediction, delta prediction, linear prediction, parallelogram prediction, etc.) to determine a predictor for the current point. A syntax element (pred_mode) signaled in geometry bitstream 203 may indicate the intra prediction strategy to use to determine the predictor for the current point.

[0196] On the other hand, if the inter flag for the current point indicates that the current point is inter predicted (“YES” branch of 800), G-PCC decoder 300 may identify a previous point in decoding order (e.g., previous point 708) (802). The previous point may have coordinates (r, phi, and laserID). G-PCC decoder 300 may then derive a quantized phi coordinate (i.e., azimuth coordinate) of the previous point (804). The quantized phi coordinate may be denoted as Q(phi). G-PCC decoder 300 may then check a reference frame (e.g., reference frame 706) for points (i.e., inter prediction points (e.g., interPredPt 704)) having quantized phi coordinates greater than the quantized phi coordinate of the

previous point (806). G-PCC decoder 300 may use the inter prediction point as a predictor for the current point (808).

[0197] Regardless of whether G-PCC decoder 300 determines the predictor for the current point using intra prediction (e.g., as described with respect to step 812) or using inter prediction (e.g., as described with respect to steps 802-808), G-PCC decoder 300 may add a delta phi multiplier (810).

[0198] The LIDAR system may scan and sample the content at a particular azimuth frequency. G-PCC encoder 200 may signal the azimuth of a point using a azimuth residual value (resAz2). The azimuth residual is not the direct difference of the azimuth of the current point and the azimuth of the predictor for the current point (which may be determined using inter or intra prediction). Rather, G-PCC encoder 200 encodes the azimuth residual as a combination of a delta phi multiplier and a residual. For example, azimSpeed may represent the azimuth difference between successive captures of one of the rotating LIDAR sensors (note there are multiple sensors in each spinning LIDAR system). Thus, if each LIDAR sensor in the spinning LIDAR system captures 1000 points per rotation, then the azimSpeed would be $(1 \ll \text{azimBitDepth}) / 1000$, effectively the azimuth difference between adjacent point captures. Here, azimBitDepth is a bitdepth used to represent the azimuth value at G-PCC encoder 200. $(1 \ll \text{azimBitDepth})$ represents one full rotation (of 360 degrees). Therefore, $(1 \ll \text{azimBitDepth}) / 1000$ represents the difference between azimuth of two captures. Although LIDAR sensors may be designed to sample at constant intervals in a rotation, the difference between azimuth values of adjacent captures may not be azimSpeed, or a multiple of azimSpeed due to noise and other inaccuracies. A first azimuth residual resAz1, which is the difference in the azimuth of the current point and the predictor, is coded as a combination of delta phi multiplier (qphi) and a second azimuth residual (resAz2). qphi may be derived as follows:

$$\begin{aligned} \text{qphi} = & \text{deltaPhi} \geq 0 ? (\text{deltaPhi} + (_geomAngularAzimuthSpeed \gg 1)) / \\ & _geomAngularAzimuthSpeed \\ & : -(\text{deltaPhi} + (_geomAngularAzimuthSpeed \gg 1)) / \\ & _geomAngularAzimuthSpeed; \end{aligned}$$

The value of coded azimuth residual, resAz2, is obtained as $\text{resAz1} - \text{deltaPhi} * \text{azimSpeed}$. Both qphi and resAz2 are signaled in the bitstream. G-PCC decoder 300 may decode resAz2 and qphi. G-PCC decoder 300 may obtain the reconstructed azimuth residual (resAz1) as $\text{resAz2} + \text{qphi} * \text{azimSpeed}$.

[0199] FIG. 9 is a conceptual diagram illustrating an example additional inter predictor point 900 obtained from the first point that has azimuth greater than an inter predictor point 914. In the inter prediction method for predictive geometry described above with respect to FIG. 7, the radius, azimuth and laserID of a current point (current point 700) are predicted based on a point (inter prediction point 704) that is near the collocated azimuth position (reference position 710) in a reference frame (reference frame 706) when inter coding is applied. In the example of FIG. 9, G-PCC encoder 200 and G-PCC decoder 300 may determine additional inter predictor point 900 using the following steps:

- a) for a given point (current point 900 of a current frame 904), determine a previous point 902 in current frame 904 (“prev dec point” in FIG. 9);
- b) determine a reference position 906 in a reference frame 908 that has the same scaled azimuth and laserID as the previous point 902 determined in step a) (“ref point with same scaled azimuth and laserID” in FIG. 9),
- c) determine a position in reference frame 908 as the first point that has an azimuth (e.g., scaled azimuth) greater than the reference position 906 determined in step b), to be used as the inter predictor point (inter pred point 910 in FIG. 9).

[0200] An additional inter predictor point may be obtained by finding the first point that has an azimuth (e.g., scaled azimuth) greater than inter prediction point 910 determined in step c) as shown in FIG. 9 (“additional inter prediction point 912” in FIG. 9). Additional signaling may be used to indicate which of the predictors is selected if inter coding has been applied. The additional inter predictor point may also be referred to as the “NextNext” inter predictor.

[0201] In some examples, G-PCC encoder 200 (e.g., arithmetic encoding unit 214 of G-PCC encoder 200) and G-PCC decoder 300 (e.g., geometry arithmetic decoding unit 302 of G-PCC decoder 300) may apply an improved context selection algorithm for coding the inter prediction flag. The inter prediction flag values of the five previously coded points may be used to select the context of the inter prediction flag in predictive geometry coding.

[0202] The aforementioned techniques may present one or more disadvantages. Inter prediction uses points in a reference frame to predict points in a given frame. In some examples, a zero compensated reference frame may be used, and in some cases a global motion compensated reference frame may be used for prediction. In the zero-motion compensated frame, the points in the reference frame are used without applying motion

compensation; the global motion compensated reference frame is obtained from the (zero compensated) reference frame by motion compensation based on a set of motion parameters; these parameters may be derived e.g., using motion estimation techniques or obtained from external means, e.g., using GPS parameters or from motion of the automobile capturing the point cloud). When motion compensation is applied to points in the reference frame, the locations of the points are modified, and this results in reduction in coding efficiency.

[0203] In accordance with one or more techniques of this disclosure, a G-PCC coder may perform resampling on a reference frame (e.g., a zero compensated reference frame, an uncompensated reference frame, etc.). Zero compensated reference typically refers to compensation with zero motion and zero rotation. In many cases, the zero compensated and uncompensated references are equivalent.

[0204] A G-PCC coder may selectively perform a resampling operation on point of a reference frame as follows:

- For each point pt in the reference frame (zero compensated), the G-PCC coder may determine whether resampling is to be applied for this point.
- If resampling is not applied, the G-PCC coder may retain the point as is.
- Otherwise (resampling is applied), the G-PCC coder may choose (e.g., select) a first set of points in the global compensated reference frame (GlobRef) that are associated with pt .
- For one or more components of pt , the G-PCC coder may derive a value from the first set of points; and use the derived value for pt for the corresponding component.
 - o E.g., if point pt was represented with (r, phi, laserId) as the components and r_d was derived radius from the first set of points, then the point pt is updated as (r_d , phi, laserId) and retained.

[0205] Individual steps of an example resampling operation that are associated with point pt are described below:

[0206] The G-PCC coder determines if resampling is to be applied to a point:

- The G-PCC coder may determine if resampling is to be applied to a point using the following condition:
 - o Whether point belongs to the road or whether the point belongs to buildings/other objects in the scene. “Road” points may be determined by

one or more thresholds that are signaled in the bitstream; for example, if two thresholds a_1 and a_2 are two z -thresholds signaled that are indicative of road points, then for points that have z value in the range of $[a_1, a_2)$, resampling is not applied; for other points resampling is applied.

- In some examples, all points may be resampled.

[0207] The G-PCC coder chooses a first set of points associated with a current point pt in the zero compensated frame, as follows:

- For a given point pt with components (r, ϕ, lId) , the first set of points in GlobRef may be chosen as points that are near the collocated point in GlobRef.
 - o E.g., points in GlobRef that have laser ID value equal to lId and azimuth in the range of $[\phi - \delta, \phi + \delta]$ where δ is fixed value based on the laser ID.
 - o In another example, an azimuth scale value may be used to determine a quantized azimuth value of all points (e.g., ϕ may be quantized to ϕ_Q). Points in a reference frame associated with a quantized azimuth value ϕ_Q may be considered be part of a bin associated with ϕ_Q . The first set of points may be determined as points in GlobRef that have laser ID equal to lId , and that belong to:
 - the bin associated with ϕ_Q ,
 - the first non-empty bin that is associated with a quantized azimuth value larger than ϕ_Q , and
 - the first non-empty bin that is associated with a quantized azimuth value smaller than ϕ_Q .
- The G-PCC coder trims the first set of points may to one or two points that have azimuth value that are closest to ϕ .
 - o In one example, if a point $P_1 (r_a, \phi_a, lId)$ belonging to GlobRef and associated with bin ϕ_Q is present:
 - If $\phi_a = \phi$, then only P_1 is retained in the first set.
 - If $\phi_a > \phi$, then P_1 is retained and another point $P_2 (r_b, \phi_b, lId)$ is chosen such that $\phi_b < \phi$.
 - If $\phi_a < \phi$, then P_1 is retained and another point $P_3 (r_c, \phi_c, lId)$ is chosen such that $\phi_c > \phi$.

If such a point P1 is not present, then two points P2 (r_b , ϕ_b , lId) and P3 (r_c , ϕ_c , lId) are chosen from the first set such that $\phi_b < \phi$ and $\phi_c > \phi$.

If one or more of points P2 or P3 are not available in GlobRef, then the unavailable point(s) may not be included in the first set of points.

[0208] The G-PCC coder may derive the radius value from points in the first set as follows (this could be applied to other components too):

- If only one point (r_x , ϕ_x , lId) is present in the first set, then the radius of point pt is updated as r_x . In other words, point pt is updated from (r , ϕ , lId) to (r_x , ϕ , lId).
- If two or more points are present in the first set, then the radius of the point pt is obtained as a value that is interpolated from the radius values of the points in the first set.

[0209] The following example of the techniques of this disclosure is provided for reference. In one example, a G-PCC coder may derive a global compensated reference frame GlobRef and update points in the zero compensated reference frame ZeroRef using resampling based on points in GlobRef. The G-PCC coder may use the updated (resampled) zero compensated reference frame for inter prediction.

[0210] FIG. 10 is a flowchart illustrating an example technique for how resampled points may be used for inter prediction, in accordance with one or more techniques of this disclosure. In the example of FIG. 10, the G-PCC coder (e.g., G-PCC encoder 200 or G-PCC decoder 300) determines that a current point in a current frame is to be predicted using an inter prediction mode. Responsive to determining to predict the current point in the current frame using the inter prediction mode, the G-PCC coder identifies a zero compensated reference frame, e.g., ZeroRef, associated with the current frame (1002). The G-PCC coder deriving a motion compensated reference frame, e.g., GlobRef, associated with the current frame (1004) and selectively resamples points in the zero compensated reference frame based on one or more points in the motion compensated reference frame to generate a resampled zero compensated reference frame (1006). The G-PCC coder inter predicts the current point in the current frame based on the resampled zero compensated reference frame (1008).

[0211] The above operation (resampling) may also be implemented as resampling the points in GlobRef using azimuth values from ZeroRef. For each point PT in ZeroRef, the

G-PCC coder may determine whether resampling is to be applied to the point. If yes, the G-PCC coder may choose a first set of points from GlobRef; obtain an interpolated radius value r_{int} from the first set of points, and update the radius value of PT with the interpolated radius value r_{int} .

[0212] FIG. 11 is a flowchart illustrating an example technique for how points in GlobRef may be resampled using points from ZeroRef, in accordance with one or more techniques of this disclosure. To selectively resample the points in the zero compensated reference frame, the G-PCC coder selects a particular point in the zero compensated reference frame (1102) and determines whether or not to resample the particular point (1104). Responsive to determining to resample the particular point, the G-PCC coder determines a first set of points in the motion compensated reference frame that are associated with the particular point (1106) and derives, based on the first set of points, a derived radius value (1108). The G-PCC coder updates, based on the derived radius value, a radius value of the particular point (1110).

[0213] FIGS. 12A-12D are conceptual diagrams illustrating how a first set of points 1202, 1204, and 1206 in GlobRef may be chosen for a point 1208 in ZeroRef, in accordance with one or more aspects of this disclosure. All points shown in FIGS. 12A-12D may have the same laser ID. Solid lines 1210, 1212, 1214, and 1216 may indicate boundaries of azimuth bins (all azimuth values within a bin are quantized to the same value). Dotted line 1218 indicates the azimuth value equal to the azimuth of point 1208. Based on the availability of point in GlobRef in the same azimuth bin as point 1208, and the difference between azimuth values of point 1202, 1204, and 1206 (if available) and point 1208, the G-PCC coder may choose the first set of points from the points 1202, 1204, and 1206.

[0214] In the example of FIG. 12A, point 1204 is present in the same azimuth bin as point 1208, and the azimuth of 1204 is equal to the azimuth of point 1208. Based on this, the G-PCC coder may select the first set of points as point 1204 (e.g., include point 1204 in the first set of points).

[0215] In the example of FIG. 12B, point 1204 is present in the same azimuth bin as PT, and the azimuth of point 1204 is greater than the azimuth of point 1208. Based on this, the G-PCC coder may select the first set of points as points 1202 and 1204 (e.g., include point 1202 and point 1204 in the first set of points).

[0216] In the example of FIG. 12C, point 1204 is present in the same azimuth bin as point 1208, and the azimuth of point 1204 is less than the azimuth of point 1208. Based on this,

the G-PCC coder may select the first set of points as 1204 and 1206 (e.g., include points 1204 and 1206 in the first set of points).

[0217] In the example of FIG. 12D, point 1204 is not present in the same azimuth bin as point 1208. Based on this, the G-PCC coder may select the first set of points as points 1202 and 1206 (e.g., include points 1202 and 1206 in the first set of points).

[0218] When the first set of points contains only one point (r_x , ϕ_x , lId), then the radius of point PT may be set equal to r_x . When the first set of points contains two points (r_x , ϕ_x , lId) and (r_y , ϕ_y , lId), the radius of point PT may be set equal to r_{xy} that may be derived as follows:

$$\begin{aligned} \text{delAzim} &= \phi_x - \phi_y \\ \text{delRad} &= r_x - r_y \\ \text{if } (\text{delRad} &= 0 \parallel \text{delAzim} = 0) \\ & \quad r_{xy} = r_x \\ \text{else} \\ & \quad r_{xy} = r_x + (\phi_{PT} - \phi_x) * \text{delRad} / \text{delAzim} \end{aligned}$$

[0219] The above example shows a linear interpolation of the radius points (e.g., an interpolation of the points of the first set of points), but higher order interpolators may also be used (e.g., cubic, etc.) and the number of points in the first set may also be more than two.

[0220] In some examples, such as when delRad is greater than a particular threshold (e.g., 1024 or 2048), the G-PCC coder may not perform interpolation.

[0221] In some examples, such as when delAzim is greater than a particular threshold, or when a quantized value of delAzim is greater than a particular threshold (delAzim quantized with azimuth scale value is more than N , where N is a fixed number, e.g. 4), the G-PCC coder may not perform interpolation.

[0222] In some examples, such as when interpolation is not performed (e.g., because threshold), the G-PCC coder may not update the point in ZeroRef. In some examples, such as when interpolation is not performed (e.g., because threshold), the G-PCC coder may set the point as to the radius of the points that has a closer azimuth to ϕ_{PT} (e.g., if $|\phi_{PT} - \phi_x| < |\phi_{PT} - \phi_y|$, then r_{xy} may be set equal to r_x , otherwise r_{xy} is set equal to r_y).

[0223] FIG. 14 is a flowchart illustrating an example technique for performing interpolation, in accordance with one or more techniques of this disclosure. In the example of FIG. 10, the G-PCC coder (e.g., G-PCC encoder 200 or G-PCC decoder 300) identifies

a reference frame (1302) and derives one or more motion compensated reference points from the reference frame (1304).

[0224] The G-PCC coder selectively resamples points in the reference frame based on the one or more motion compensated reference points to generate a resampled reference frame (1306). To selectively resample the points in the reference frame, G-PCC coder may be configured to select a particular point in the reference frame; determine whether or not to resample the particular point; and responsive to determining to resample the particular point, determine a first set of points in the reference frame that are associated with the particular point, wherein the first set of points includes at least one of the one or more motion compensated reference points, derive, based on the first set of points, a derived radius value, and update, based on the derived radius value, a radius value of the particular point. To derive the derived radius value based on the first set of points, the G-PCC coder may be configured to selectively interpolate radius values for a plurality of points to determine the derived radius value, wherein the first set of points includes the plurality of points. To selectively interpolate the radius values for the plurality of points, G-PCC coder may be configured to determine not to interpolate radius values for a first point and a second point of the plurality of points in response to one or both of a difference of radius values of the first point and second point being greater than a first threshold value or a difference of azimuth values of the first point and the second point being greater than a second threshold value.

The G-PCC coder may be configured to refrain from updating the radius value of the particular point based on the derived radius value in response to determining to not interpolate the radius values for the plurality of points. The G-PCC coder may be configured to set the radius value of the particular point as a radius value of a single point of the first set of points in response to determining to not interpolate the radius values of the plurality of points. The single point of the first set of points may, for example, be the point that has an azimuth value closest to an azimuth value of the particular point.

[0225] The G-PCC coder inter predicts the current point in the current frame based on the resampled reference frame (1308). The G-PCC coder may reconstruct a point cloud that includes the current point.

[0226] FIG. 14 is a conceptual diagram illustrating an example range-finding system 1400 that may be used with one or more techniques of this disclosure. In the example of FIG. 14, range-finding system 1400 includes an illuminator 1402 and a sensor 1404.

Illuminator 1402 may emit light 1406. In some examples, illuminator 1402 may emit light 1406 as one or more laser beams. Light 1406 may be in one or more wavelengths, such as an infrared wavelength or a visible light wavelength. In other examples, light 1406 is not a coherent, laser light. When light 1406 encounters an object, such as object 1408, light 1406 creates returning light 1410. Returning light 1410 may include backscattered and/or reflected light. Returning light 1410 may pass through a lens 1411 that directs returning light 1410 to create an image 1412 of object 1408 on sensor 1404. Sensor 1404 generates signals 1414 based on image 1412. Image 1412 may comprise a set of points (e.g., as represented by dots in image 1412 of FIG. 14).

[0227] In some examples, illuminator 1402 and sensor 1404 may be mounted on a spinning structure so that illuminator 1402 and sensor 1404 capture a 360A-degree view of an environment. In other examples, range-finding system 1400 may include one or more optical components (e.g., mirrors, collimators, diffraction gratings, etc.) that enable illuminator 1402 and sensor 1404 to detect ranges of objects within a specific range (e.g., up to 360A-degrees). Although the example of FIG. 14 only shows a single illuminator 1402 and sensor 1404, range-finding system 1400 may include multiple sets of illuminators and sensors.

[0228] In some examples, illuminator 1402 generates a structured light pattern. In such examples, range-finding system 1400 may include multiple sensors 1404 upon which respective images of the structured light pattern are formed. Range-finding system 1400 may use disparities between the images of the structured light pattern to determine a distance to an object 1408 from which the structured light pattern backscatters. Structured light-based range-finding systems may have a high level of accuracy (e.g., accuracy in the sub-millimeter range), when object 1408 is relatively close to sensor 1404 (e.g., 0.2 meters to 2 meters). This high level of accuracy may be useful in facial recognition applications, such as unlocking mobile devices (e.g., mobile phones, tablet computers, etc.) and for security applications.

[0229] In some examples, range-finding system 1400 is a time of flight (ToF)-based system. In some examples where range-finding system 1400 is a ToF-based system, illuminator 1402 generates pulses of light. In other words, illuminator 1402 may modulate the amplitude of emitted light 1406. In such examples, sensor 1404 detects returning light 1410 from the pulses of light 1406 generated by illuminator 1402. Range-finding system 1400 may then determine a distance to object 1408 from which light 1406 backscatters based on a delay between when light 1406 was emitted and detected and the

known speed of light in air). In some examples, rather than (or in addition to) modulating the amplitude of the emitted light 1406, illuminator 1402 may modulate the phase of the emitted light 1406. In such examples, sensor 1404 may detect the phase of returning light 1410 from object 1408 and determine distances to points on object 1408 using the speed of light and based on time differences between when illuminator 1402 generated light 1406 at a specific phase and when sensor 1404 detected returning light 1410 at the specific phase.

[0230] In other examples, a point cloud may be generated without using illuminator 1402. For instance, in some examples, sensors 1404 of range-finding system 1400 may include two or more optical cameras. In such examples, range-finding system 1400 may use the optical cameras to capture stereo images of the environment, including object 1408. Range-finding system 1400 may include a point cloud generator 1416 that may calculate the disparities between locations in the stereo images. Range-finding system 1400 may then use the disparities to determine distances to the locations shown in the stereo images. From these distances, point cloud generator 1416 may generate a point cloud.

[0231] Sensors 1404 may also detect other attributes of object 1408, such as color and reflectance information. In the example of FIG. 14, a point cloud generator 1416 may generate a point cloud based on signals 1414 generated by sensor 1404. Range-finding system 1400 and/or point cloud generator 1416 may form part of data source 104 (FIG. 1). Hence, a point cloud generated by range-finding system 1400 may be encoded and/or decoded according to any of the techniques of this disclosure.

[0232] FIG. 15 is a conceptual diagram illustrating an example vehicle-based scenario in which one or more techniques of this disclosure may be used. In the example of FIG. 15, a vehicle 1500 includes a range-finding system 1502. Range-finding system 1502 may be implemented in the manner discussed with respect to FIG. 15. Although not shown in the example of FIG. 15, vehicle 1500 may also include a data source, such as data source 104 (FIG. 1), and a G-PCC encoder, such as G-PCC encoder 200 (FIG. 1). In the example of FIG. 15, range-finding system 1502 emits laser beams 1504 that reflect off pedestrians 1506A or other objects in a roadway. The data source of vehicle 1500 may generate a point cloud based on signals generated by range-finding system 1502. The G-PCC encoder of vehicle 1500 may encode the point cloud to generate bitstreams 1508, such as geometry bitstream (FIG. 2) and attribute bitstream (FIG. 2). Bitstreams 1508 may include many fewer bits than the unencoded point cloud obtained by the G-PCC encoder. In some examples, the G-PCC encoder of vehicle 1500 may encode the bitstreams 1508

using radius interpolation as described above. In some examples, the G-PCC decoder of vehicle 1510 may decode the bitstreams 1508 using radius interpolation as described above.

[0233] An output interface of vehicle 1500 (e.g., output interface 108 (FIG. 1)) may transmit bitstreams 1508 to one or more other devices. Bitstreams 1508 may include many fewer bits than the unencoded point cloud obtained by the G-PCC encoder. Thus, vehicle 1500 may be able to transmit bitstreams 1508 to other devices more quickly than the unencoded point cloud data. Additionally, bitstreams 1508 may require less data storage capacity.

[0234] In the example of FIG. 15, vehicle 1500 may transmit bitstreams 1508 to another vehicle 1510. Vehicle 1510 may include a G-PCC decoder, such as G-PCC decoder 300 (FIG. 1). The G-PCC decoder of vehicle 1510 may decode bitstreams 1508 to reconstruct the point cloud. Vehicle 1510 may use the reconstructed point cloud for various purposes. For instance, vehicle 1510 may determine based on the reconstructed point cloud that pedestrians 1506 are in the roadway ahead of vehicle 1500 and therefore start slowing down, e.g., even before a driver of vehicle 1510 realizes that pedestrians 1206 are in the roadway. Thus, in some examples, vehicle 1510 may perform an autonomous navigation operation based on the reconstructed point cloud.

[0235] Additionally, or alternatively, vehicle 1500 may transmit bitstreams 1508 to a server system 1512. Server system 1512 may use bitstreams 1508 for various purposes. For example, server system 1512 may store bitstreams 1508 for subsequent reconstruction of the point clouds. In this example, server system 1512 may use the point clouds along with other data (e.g., vehicle telemetry data generated by vehicle 1500) to train an autonomous driving system. In other example, server system 1512 may store bitstreams 1508 for subsequent reconstruction for forensic crash investigations.

[0236] FIG. 16 is a conceptual diagram illustrating an example extended reality system in which one or more techniques of this disclosure may be used. Extended reality (XR) is a term used to cover a range of technologies that includes augmented reality (AR), mixed reality (MR), and virtual reality (VR). In the example of FIG. 16, a user 1600 is located in a first location 1602. User 1600 wears an XR headset 1604. As an alternative to XR headset 1604, user 1600 may use a mobile device (e.g., mobile phone, tablet computer, etc.). XR headset 1604 includes a depth detection sensor, such as a range-finding system, that detects positions of points on objects 1606 at location 1602. A data source of XR headset 1604 may use the signals generated by the depth detection sensor

to generate a point cloud representation of objects 1606 at location 1602. XR headset 1604 may include a G-PCC encoder (e.g., G-PCC encoder 200 of FIG. 1) that is configured to encode the point cloud to generate bitstreams 1608. In some examples, the G-PCC encoder of XR headset 1604 may use radius interpolation when encoding the point cloud, as described above.

[0237] XR headset 1604 may transmit bitstreams 1608 (e.g., via a network such as the Internet) to an XR headset 1610 worn by a user 1612 at a second location 1614. XR headset 1610 may decode bitstreams 1608 to reconstruct the point cloud. In some examples, the G-PCC decoder of XR headset 1610 may use radius interpolation when decoding the point cloud, as described above.

[0238] XR headset 1610 may use the point cloud to generate an XR visualization (e.g., an AR, MR, VR visualization) representing objects 1606 at location 1602. Thus, in some examples, such as when XR headset 1610 generates an VR visualization, user 1612 may have a 3D immersive experience of location 1602. In some examples, XR headset 1610 may determine a position of a virtual object based on the reconstructed point cloud. For instance, XR headset 1610 may determine, based on the reconstructed point cloud, that an environment (e.g., location 1602) includes a flat surface and then determine that a virtual object (e.g., a cartoon character) is to be positioned on the flat surface. XR headset 1610 may generate an XR visualization in which the virtual object is at the determined position. For instance, XR headset 1610 may show the cartoon character sitting on the flat surface.

[0239] FIG. 17 is a conceptual diagram illustrating an example mobile device system in which one or more techniques of this disclosure may be used. In the example of FIG. 17, a mobile device 1700, such as a mobile phone or tablet computer, includes a range-finding system, such as a LIDAR system, that detects positions of points on objects 1702 in an environment of mobile device 1700. A data source of mobile device 1700 may use the signals generated by the depth detection sensor to generate a point cloud representation of objects 1702. Mobile device 1700 may include a G-PCC encoder (e.g., G-PCC encoder 200 of FIG. 1) that is configured to encode the point cloud to generate bitstreams 1704.

[0240] In the example of FIG. 17, mobile device 1700 may transmit bitstreams to a remote device 1706, such as a server system or other mobile device. Remote device 1706 may decode bitstreams 1704 to reconstruct the point cloud. In some examples, the G-PCC decoder of remote device 1706 may use radius interpolation when decoding the point cloud, as described above.

[0241] Remote device 1706 may use the point cloud for various purposes. For example, remote device 1706 may use the point cloud to generate a map of environment of mobile device 1700. For instance, remote device 1706 may generate a map of an interior of a building based on the reconstructed point cloud. In another example, remote device 1706 may generate imagery (e.g., computer graphics) based on the point cloud. For instance, remote device 1706 may use points of the point cloud as vertices of polygons and use color attributes of the points as the basis for shading the polygons. In some examples, remote device 1706 may use the reconstructed point cloud for facial recognition or other security applications.

[0242] The following numbered clauses may illustrate one or more aspects of the disclosure:

[0243] Clause 1A: A method of processing a point cloud, the method comprising: responsive to determining to predict a current point in a current frame using an inter prediction mode: identifying a zero compensated reference frame; deriving a motion compensated reference frame; selectively resampling points in the zero compensated reference frame based on one or more points in the motion compensated reference frame to generate a resampled zero compensated reference frame; and inter predict the current point in the current frame based on the resampled zero compensated reference frame.

[0244] Clause 2A: The method of clause 1A, wherein the zero compensated reference frame is ZeroRef, and the motion compensated reference frame is GlobRef.

[0245] Clause 3A: The method of clause 1A or clause 2A, wherein selectively resampling the points in the zero compensated reference frame comprises: selecting a particular point in the zero compensated reference frame; determining whether or not to resample the particular point; responsive to determining to resample the particular point: determining a first set of points in the motion compensated reference frame that are associated with the particular point; deriving, based on the first set of points, a derived radius value; and update, based on the derived radius value, a radius value of the particular point.

[0246] Clause 4A: The method of clause 3A, wherein deriving the derived radius value based on the first set of points comprises: selectively interpolating, where the first set of points includes a plurality of points, the derived radius value based on the plurality of points.

[0247] Clause 5A: The method of clause 4A, wherein selectively interpolating the derived radius value comprises: determining not to interpolate the derived radius value

where a difference of radius values of the plurality of points is greater than a threshold value.

[0248] Clause 6A: The method of clause 4A, wherein selectively interpolating the derived radius value comprises: determining not to interpolate the derived radius value where a difference of azimuth values of the plurality of points is greater than a threshold value.

[0249] Clause 7A: The method of any of clauses 4A-6A, further comprising: refraining, responsive to determining to not interpolate the derived radius value, from updating the radius value of the particular point based on the derived radius value.

[0250] Clause 8A: The method of any of clauses 4A-6A, further comprising: set, responsive to determining to not interpolate the derived radius value, the radius value of the particular point as a radius value of a single point of the first set of points.

[0251] Clause 9A: The method of clause 8A, wherein the single point of the first set of points has an azimuth value closest to an azimuth value of the particular point.

[0252] Clause 10A: A device for processing a point cloud, the device comprising one or more means for performing the method of any of clauses 1A-9A.

[0253] Clause 11A: The device of clause 10A, wherein the one or more means comprise one or more processors implemented in circuitry.

[0254] Clause 12A: The device of any of clauses 10A or 11A, further comprising a memory to store the data representing the point cloud. Clause 13A: The device of any of clauses 10A-12A, wherein the device comprises a decoder.

[0255] Clause 14A: The device of any of clauses 10A-13A, wherein the device comprises an encoder.

[0256] Clause 15A: The device of any of clauses 10A-14A, further comprising a device to generate the point cloud.

[0257] Clause 16A: The device of any of clauses 10A-15A, further comprising a display to present imagery based on the point cloud.

[0258] Clause 17A: A computer-readable storage medium having stored thereon instructions that, when executed, cause one or more processors to perform the method of any of clauses 1A-9A.

[0259] Clause 1B: A device for processing point cloud data, the device comprising: a memory configured to store the point cloud data; and one or more processors, implemented in circuitry, and configured to, responsive to a determination to predict a current point in a current frame using an inter prediction mode: identify a reference frame;

derive one or more motion compensated reference points from the reference frame; selectively resample points in the reference frame based on the one or more motion compensated reference points to generate a resampled reference frame; and inter predict the current point in the current frame based on the resampled reference frame.

[0260] Clause 2B: The device of clause 1B, wherein to selectively resample the points in the reference frame, the one or more processors are further configured to: select a particular point in the reference frame; determine whether or not to resample the particular point; responsive to determining to resample the particular point: determine a first set of points in the reference frame that are associated with the particular point, wherein the first set of points includes at least one of the one or more motion compensated reference points; derive, based on the first set of points, a derived radius value; and update, based on the derived radius value, a radius value of the particular point.

[0261] Clause 3B: The device of clause 2B, wherein to derive the derived radius value based on the first set of points, the one or more processors are further configured to: selectively interpolate radius values for a plurality of points to determine the derived radius value, wherein the first set of points includes the plurality of points.

[0262] Clause 4B: The device of clause 3B, wherein to selectively interpolate the radius values for the plurality of points, the one or more processors are further configured to: determine not to interpolate radius values for a first point and a second point of the plurality of points in response to one or both of a difference of radius values of the first point and second point being greater than a first threshold value or a difference of azimuth values of the first point and the second point being greater than a second threshold value.

[0263] Clause 5B: The device of clause 4B, wherein the one or more processors are further configured to: refrain from updating the radius value of the particular point based on the derived radius value in response to a determination to not interpolate the radius values for the plurality of points.

[0264] Clause 6B: The device of clause 4B, wherein the one or more processors are further configured to: set the radius value of the particular point as a radius value of a single point of the first set of points in response to a determination to not interpolate the radius values for the plurality of points.

[0265] Clause 7B: The device of clause 6B, wherein the single point of the first set of points has an azimuth value closest to an azimuth value of the particular point.

[0266] Clause 8B: The device of any of clauses 1B-7B, wherein the one or more processors are further configured to reconstruct a point cloud that includes the current point to form a reconstructed point cloud.

[0267] Clause 9B: The device of clause 8B, wherein the one or more processors are further configured to generate a map of an interior of a building based on the reconstructed point cloud.

[0268] Clause 10B: The device of clause 8B or 9B, wherein the one or more processors are further configured to perform an autonomous navigation operation based on the reconstructed point cloud.

[0269] Clause 11B: The device of any of clauses 8B-10B, wherein the one or more processors are further configured to generate computer graphics based on the reconstructed point cloud.

[0270] Clause 12B: The device of any of clauses 8B-11B, wherein the one or more processors are configured to: determine a position of a virtual object based on the reconstructed point cloud; and generate an extended reality (XR) visualization in which the virtual object is at the determined position.

[0271] Clause 13B: The device of any of clauses 8B-12B, further comprising a display to present imagery based on the reconstructed point cloud.

[0272] Clause 14B: The device of any of clauses 1B-13B, wherein the device comprises one or more of a mobile phone, a tablet computer, a vehicle, or an extended reality device.

[0273] Clause 15B: A method of processing a point cloud, the method comprising: responsive to determining to predict a current point in a current frame using an inter prediction mode: identifying a reference frame; deriving one or more motion compensated reference points from the reference frame; selectively resampling points in the reference frame based on the one or more points motion compensated reference points to generate a resampled reference frame; and inter predicting the current point in the current frame based on the resampled reference frame.

[0274] Clause 16B: The method of clause 15B, wherein selectively resampling the points in the reference frame comprises: selecting a particular point in the reference frame; determining whether or not to resample the particular point; responsive to determining to resample the particular point: determining a first set of points in the reference frame that are associated with the particular point, wherein the first set of points includes at least one of the one or more motion compensated reference points; deriving, based on the first set

of points, a derived radius value; and updating, based on the derived radius value, a radius value of the particular point.

[0275] Clause 17B: The method of clause 16B, wherein deriving the derived radius value based on the first set of points comprises: selectively interpolating radius values for a plurality of points to determine the derived radius value, wherein the first set of points includes the plurality of points.

[0276] Clause 18B: The method of clause 17B, wherein selectively interpolating the derived radius value comprises: determining not to interpolate radius values for a first point and a second point of the plurality of points in response to one or both of a difference of radius values of the first point and second point being greater than a first threshold value or a difference of azimuth values of the first point and the second point being greater than a second threshold value.

[0277] Clause 19B: The method of clause 18B, further comprising: refraining from updating the radius value of the particular point based on the derived radius value in response to determining to not interpolate the radius values for the plurality of points.

[0278] Clause 20B: The method of clause 18B, further comprising: setting the radius value of the particular point as a radius value of a single point of the first set of points in response to determining to not interpolate the radius values of the plurality of points.

[0279] Clause 21B: The method of clause 20B, wherein the single point of the first set of points has an azimuth value closest to an azimuth value of the particular point.

[0280] Clause 22B: The method of any of clauses 15B-21B, wherein the method of processing the point cloud comprises a method of encoding the point cloud.

[0281] Clause 23B: The method of any of clauses 15B-21B, wherein the method of processing the point cloud comprises a method of decoding the point cloud.

[0282] Clause 24B: A computer-readable storage medium storing instructions that when executed by one or more processors cause the one or more processors to: responsive to determining to predict a current point in a current frame using an inter prediction mode: identify a reference frame; derive one or more motion compensated reference points from the reference frame; selectively resample points in the reference frame based on the one or more points motion compensated reference points to generate a resampled reference frame; and inter predict the current point in the current frame based on the resampled reference frame.

[0283] Clause 25B: The computer-readable storage medium of clause 24B, wherein to selectively resampling the points in the reference frame, the instructions cause the one or

more processors to: select a particular point in the reference frame; determine whether or not to resample the particular point; responsive to determining to resample the particular point: determine a first set of points in the reference frame that are associated with the particular point, wherein the first set of points includes at least one of the one or more motion compensated reference points; derive, based on the first set of points, a derived radius value; and update, based on the derived radius value, a radius value of the particular point.

[0284] Clause 26B: The computer-readable storage medium of clause 25B, wherein to derive the derived radius value based on the first set of points, the instructions cause the one or more processors to: selectively interpolate radius values for a plurality of points to determine the derived radius value, wherein the first set of points includes the plurality of points.

[0285] Clause 27B: The computer-readable storage medium of clause 26B, wherein to selectively interpolating the derived radius value, the instructions cause the one or more processors to: determine not to interpolate radius values for a first point and a second point of the plurality of points in response to one or both of a difference of radius values of the first point and second point being greater than a first threshold value or a difference of azimuth values of the first point and the second point being greater than a second threshold value.

[0286] Clause 28B: The computer-readable storage medium of clause 27B, wherein the instructions further cause the one or more processors to: refrain from updating the radius value of the particular point based on the derived radius value in response to determining to not interpolate the radius values for the plurality of points.

[0287] Clause 29B: The computer-readable storage medium of clause 27B, wherein the instructions further cause the one or more processors to: set the radius value of the particular point as a radius value of a single point of the first set of points in response to determining to not interpolate the radius values of the plurality of points.

[0288] Clause 30B: The computer-readable storage medium of clause 29B, wherein the single point of the first set of points has an azimuth value closest to an azimuth value of the particular point.

[0289] It is to be recognized that depending on the example, certain acts or events of any of the techniques described herein can be performed in a different sequence, may be added, merged, or left out altogether (e.g., not all described acts or events are necessary for the practice of the techniques). Moreover, in certain examples, acts or events may be

performed concurrently, e.g., through multi-threaded processing, interrupt processing, or multiple processors, rather than sequentially.

[0290] In one or more examples, the functions described may be implemented in hardware, software, firmware, or any combination thereof. If implemented in software, the functions may be stored on or transmitted over as one or more instructions or code on a computer-readable medium and executed by a hardware-based processing unit. Computer-readable media may include computer-readable storage media, which corresponds to a tangible medium such as data storage media, or communication media including any medium that facilitates transfer of a computer program from one place to another, e.g., according to a communication protocol. In this manner, computer-readable media generally may correspond to (1) tangible computer-readable storage media which is non-transitory or (2) a communication medium such as a signal or carrier wave. Data storage media may be any available media that can be accessed by one or more computers or one or more processors to retrieve instructions, code and/or data structures for implementation of the techniques described in this disclosure. A computer program product may include a computer-readable medium.

[0291] By way of example, and not limitation, such computer-readable storage media can comprise RAM, ROM, EEPROM, CD-ROM or other optical disk storage, magnetic disk storage, or other magnetic storage devices, flash memory, or any other medium that can be used to store desired program code in the form of instructions or data structures and that can be accessed by a computer. Also, any connection is properly termed a computer-readable medium. For example, if instructions are transmitted from a website, server, or other remote source using a coaxial cable, fiber optic cable, twisted pair, digital subscriber line (DSL), or wireless technologies such as infrared, radio, and microwave, then the coaxial cable, fiber optic cable, twisted pair, DSL, or wireless technologies such as infrared, radio, and microwave are included in the definition of medium. It should be understood, however, that computer-readable storage media and data storage media do not include connections, carrier waves, signals, or other transitory media, but are instead directed to non-transitory, tangible storage media. Disk and disc, as used herein, includes compact disc (CD), laser disc, optical disc, digital versatile disc (DVD), floppy disk and Blu-ray disc, where disks usually reproduce data magnetically, while discs reproduce data optically with lasers. Combinations of the above should also be included within the scope of computer-readable media.

[0292] Instructions may be executed by one or more processors, such as one or more digital signal processors (DSPs), general purpose microprocessors, application specific integrated circuits (ASICs), field programmable gate arrays (FPGAs), or other equivalent integrated or discrete logic circuitry. Accordingly, the terms “processor” and “processing circuitry,” as used herein may refer to any of the foregoing structures or any other structure suitable for implementation of the techniques described herein. In addition, in some aspects, the functionality described herein may be provided within dedicated hardware and/or software modules configured for encoding and decoding, or incorporated in a combined codec. Also, the techniques could be fully implemented in one or more circuits or logic elements.

[0293] The techniques of this disclosure may be implemented in a wide variety of devices or apparatuses, including a wireless handset, an integrated circuit (IC) or a set of ICs (e.g., a chip set). Various components, modules, or units are described in this disclosure to emphasize functional aspects of devices configured to perform the disclosed techniques, but do not necessarily require realization by different hardware units. Rather, as described above, various units may be combined in a codec hardware unit or provided by a collection of interoperative hardware units, including one or more processors as described above, in conjunction with suitable software and/or firmware.

[0294] Various examples have been described. These and other examples are within the scope of the following clauses.

WHAT IS CLAIMED IS:

1. A device for processing point cloud data, the device comprising:
 - a memory configured to store the point cloud data; and
 - one or more processors, implemented in circuitry, and configured to, responsive to a determination to predict a current point in a current frame using an inter prediction mode:
 - identify a reference frame;
 - derive one or more motion compensated reference points from the reference frame;
 - selectively resample points in the reference frame based on the one or more motion compensated reference points to generate a resampled reference frame;
 - and
 - inter predict the current point in the current frame based on the resampled reference frame.
2. The device of claim 1, wherein to selectively resample the points in the reference frame, the one or more processors are further configured to:
 - select a particular point in the reference frame;
 - determine whether or not to resample the particular point;
 - responsive to a determination to resample the particular point:
 - determine a first set of points in the reference frame that are associated with the particular point, wherein the first set of points includes at least one of the one or more motion compensated reference points;
 - derive, based on the first set of points, a derived radius value; and
 - update, based on the derived radius value, a radius value of the particular point.
3. The device of claim 2, wherein to derive the derived radius value based on the first set of points, the one or more processors are further configured to:
 - selectively interpolate radius values for a plurality of points to determine the derived radius value, wherein the first set of points includes the plurality of points.

4. The device of claim 3, wherein to selectively interpolate the radius values for the plurality of points, the one or more processors are further configured to:

determine not to interpolate radius values for a first point and a second point of the plurality of points in response to one or both of a difference of radius values of the first point and second point being greater than a first threshold value or a difference of azimuth values of the first point and the second point being greater than a second threshold value.

5. The device of claim 4, wherein the one or more processors are further configured to:

refrain from updating the radius value of the particular point based on the derived radius value in response to a determination to not interpolate the radius values for the plurality of points.

6. The device of claim 4, wherein the one or more processors are further configured to:

set the radius value of the particular point as a radius value of a single point of the first set of points in response to a determination to not interpolate the radius values for the plurality of points.

7. The device of claim 6, wherein the single point of the first set of points has an azimuth value closest to an azimuth value of the particular point.

8. The device of claim 1, wherein the one or more processors are further configured to reconstruct a point cloud that includes the current point to form a reconstructed point cloud.

9. The device of claim 8, wherein the one or more processors are further configured to generate a map of an interior of a building based on the reconstructed point cloud.

10. The device of claim 8, wherein the one or more processors are further configured to perform an autonomous navigation operation based on the reconstructed point cloud.

11. The device of claim 8, wherein the one or more processors are further configured to generate computer graphics based on the reconstructed point cloud.
12. The device of claim 8, wherein the one or more processors are configured to:
determine a position of a virtual object based on the reconstructed point cloud;
and
generate an extended reality (XR) visualization in which the virtual object is at the determined position.
13. The device of claim 8, further comprising a display to present imagery based on the reconstructed point cloud.
14. The device of claim 1, wherein the device comprises one or more of a mobile phone, a tablet computer, a vehicle, or an extended reality device.
15. A method of processing a point cloud, the method comprising:
responsive to determining to predict a current point in a current frame using an inter prediction mode:
identifying a reference frame;
deriving one or more motion compensated reference points from the reference frame;
selectively resampling points in the reference frame based on the one or more points motion compensated reference points to generate a resampled reference frame; and
inter predicting the current point in the current frame based on the resampled reference frame.

16. The method of claim 15, wherein selectively resampling the points in the reference frame comprises:

- selecting a particular point in the reference frame;
- determining whether or not to resample the particular point;
- responsive to determining to resample the particular point:
 - determining a first set of points in the reference frame that are associated with the particular point, wherein the first set of points includes at least one of the one or more motion compensated reference points;
 - deriving, based on the first set of points, a derived radius value; and
 - updating, based on the derived radius value, a radius value of the particular point.

17. The method of claim 16, wherein deriving the derived radius value based on the first set of points comprises:

- selectively interpolating radius values for a plurality of points to determine the derived radius value, wherein the first set of points includes the plurality of points.

18. The method of claim 17, wherein selectively interpolating the derived radius value comprises:

- determining not to interpolate radius values for a first point and a second point of the plurality of points in response to one or both of a difference of radius values of the first point and second point being greater than a first threshold value or a difference of azimuth values of the first point and the second point being greater than a second threshold value.

19. The method of claim 18, further comprising:

- refraining from updating the radius value of the particular point based on the derived radius value in response to determining to not interpolate the radius values for the plurality of points.

20. The method of claim 18, further comprising:

- setting the radius value of the particular point as a radius value of a single point of the first set of points in response to determining to not interpolate the radius values of the plurality of points.

21. The method of claim 20, wherein the single point of the first set of points has an azimuth value closest to an azimuth value of the particular point.

22. The method of claim 15, wherein the method of processing the point cloud comprises a method of encoding the point cloud.

23. The method of claim 15, wherein the method of processing the point cloud comprises a method of decoding the point cloud.

24. A computer-readable storage medium storing instructions that when executed by one or more processors cause the one or more processors to:

responsive to determining to predict a current point in a current frame using an inter prediction mode:

identify a reference frame;

derive one or more motion compensated reference points from the reference frame;

selectively resample points in the reference frame based on the one or more points motion compensated reference points to generate a resampled reference frame; and

inter predict the current point in the current frame based on the resampled reference frame.

25. The computer-readable storage medium of claim 24, wherein to selectively resampling the points in the reference frame, the instructions cause the one or more processors to:

- select a particular point in the reference frame;
- determine whether or not to resample the particular point;
- responsive to determining to resample the particular point:
 - determine a first set of points in the reference frame that are associated with the particular point, wherein the first set of points includes at least one of the one or more motion compensated reference points;
 - derive, based on the first set of points, a derived radius value; and
 - update, based on the derived radius value, a radius value of the particular point.

26. The computer-readable storage medium of claim 25, wherein to derive the derived radius value based on the first set of points, the instructions cause the one or more processors to:

- selectively interpolate radius values for a plurality of points to determine the derived radius value, wherein the first set of points includes the plurality of points.

27. The computer-readable storage medium of claim 26, wherein to selectively interpolating the derived radius value, the instructions cause the one or more processors to:

- determine not to interpolate radius values for a first point and a second point of the plurality of points in response to one or both of a difference of radius values of the first point and second point being greater than a first threshold value or a difference of azimuth values of the first point and the second point being greater than a second threshold value.

28. The computer-readable storage medium of claim 27, wherein the instructions further cause the one or more processors to:

- refrain from updating the radius value of the particular point based on the derived radius value in response to determining to not interpolate the radius values for the plurality of points.

29. The computer-readable storage medium of claim 27, wherein the instructions further cause the one or more processors to:

set the radius value of the particular point as a radius value of a single point of the first set of points in response to determining to not interpolate the radius values of the plurality of points.

30. The computer-readable storage medium of claim 29, wherein the single point of the first set of points has an azimuth value closest to an azimuth value of the particular point.