



(51) International Patent Classification:
G06T 17/05 (2011.01)

(21) International Application Number:
PCT/US2011/067958

(22) International Filing Date:
29 December 2011 (29.12.2011)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
13/007,161 14 January 2011 (14.01.2011) US

(71) Applicant (for all designated States except US): **INTEL CORPORATION** [US/US]; 2200 Mission College Boulevard, Santa Clara, California 95052 (US).

(72) Inventor; and

(75) Inventor/Applicant (for US only): **SUN, Wei** [CA/US]; 80 Descanso Drive, Apartment 2308, San Jose, California 95134 (US).

(74) Agent: **TROP, Timothy N.**; Trop, Pruner & Hu, P.C., 1616 S. Voss Rd., Ste. 750, Houston, Texas 77057-2631 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO,

DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

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

Declarations under Rule 4.17:

- as to the identity of the inventor (Rule 4.17(i))
- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))
- as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))

Published:

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

(54) Title: PLANETARY SCALE OBJECT RENDERING

(57) Abstract: A planetary scale object rendering system may provide a plurality of spherical mesh topologies of different levels of detail made up of abutting faces. Tessellation of the faces may be provided at continuous levels of detail. The geometry data at different locations and levels of detail in the mesh topology may be randomly accessible. The geometric data may be organized in quadrees in some embodiments. A layered quadtree data structure may be formed from the bottom layer up with successively higher parent layers formed from the root nodes of previous child layers. The quadtree structure may be randomly accessed from the top down.



- 1 -

Planetary Scale Object Rendering

Background

[0001] This relates generally to planetary scale object rendering. As used herein, planetary scale object rendering involves rendering objects of at least 100,000 times difference in scale.

[0002] With planetary object rendering, used, for example, in Google Earth, a user may view the entire Earth and see one level of resolution and then may, effectively, zoom in to a particular region on the Earth and see a greater resolution with respect to the smaller area. Thus, Google Earth allows users to select particular regions to view, how much of those regions to view, and thereby the system determines the resolution that the user may see.

[0003] Planetary viewers may provide convenient platforms for visualizing various scientific data, as well as for personal, social, and mass media. Planetary viewers may also be used in navigation, eco-management, geographical oversight and planning, environmental science, emergency planning and response, as well as for entertainment.

[0004] Rendering a planetary scale object, such as the Earth, with high precision in real time is challenging for 3D geometry processing due to the large variations in viewing distance and viewing area and the need for real time retrieval of a very large amount of two dimensional data, including surface elevation data. For example, a planetary viewer may view regions as close as one meter and as far as tens of thousands of kilometers, with viewing areas as small as one square meter and as large as the entire hemisphere.

Brief Description Of The Drawings

[0005] Figure 1 is a perspective view of three of a plurality of mesh topologies for one embodiment;

- 2 -

[0006] Figure 2 is a flow chart for a geometry processing pipeline according to one embodiment;

[0007] Figure 3 is a depiction of patch tessellation according to one embodiment;

[0008] Figure 4 is a plot showing displacement mapping according to one embodiment;

[0009] Figure 5 is a flow chart for offline processing according to one embodiment;

[0010] Figure 6 is a layout of a quadtree hierarchy for one embodiment;

[0011] Figure 7 is a depiction of a hierarchical quadtree for one embodiment;

[0012] Figure 8 is a flow chart for a run time retrieval of topographic data according to one embodiment; and

[0013] Figure 9 is a hardware depiction for one embodiment.

Detailed Description

[0014] In accordance with some embodiments of the present invention, the planetary object is the Earth or another celestial planet, but the present invention can be applied to any planetary scale objects (scale greater than 100,000X). In one embodiment, a planet is approximated as a perfect sphere. This may simplify sampling and interpolation and is easier for calculating surface normals.

[0015] In one embodiment, there may be a plurality of fixed or predefined, discrete levels of detail spherical mesh topologies. For example, in one embodiment, eight discrete levels of detail, implemented by spherical mesh topologies, may span from a coarsest level of detail to a finest level of detail. In one embodiment, the coarsest level of detail is 1667x1667 kilometers per patch and the finest level of detail patch is 13.02x13.02 square kilometers. As used herein, a face is quadrilateral region on a spherical mesh topology (except for the polar area where a face is a triangle) that represents a patch on a real geometrical object.

- 3 -

[0016] Referring to Figure 1, three of a plurality of spherical mesh topologies 10 to implement a predetermined number of discrete levels of detail may be used in one embodiment. Only three topologies are depicted to illustrate increasing level of detail from left to right. Each mesh topology includes a plurality of faces 12 that are idealized and simplified as flat surfaces. Thus, each mesh topology is made up of a plurality of planar faces representing a given resolution. Each abutting planar face represents a roughly similarly sized and shaped spherical (i.e. non-planar) section, called a patch, of the actual spherical object.

[0017] In some embodiments, the fixed or discrete levels of detail implemented by the mesh topologies are supplemented by continuous levels of detail for the mesh topologies and topography data. In some embodiments, efficient data structures may be used for fast network transmission, fast hard disk input/output (I/O), and efficient memory caching.

[0018] Referring to Figure 2, a geometry processing pipeline for planetary scale object rendering begins, in one embodiment, by creating hierarchical mesh models, as indicated in block 14. The hierarchical mesh models are the predetermined number of discrete levels of detail, as depicted in Figure 1. While an example is given in which eight mesh models are used, the number of mesh models may be chosen based on the capabilities of the processing system.

[0019] The mesh topologies are not stored, in some embodiments, because they likely would be of a size that would not fit in local storage. It would take two 32-bit floating point values to store the geographical position (i.e. longitude and latitude) of each vertex, one 32-bit integer to index each vertex, one 32-bit integer to index each edge, one 32-bit integer to index each face, and extra index duplications to store the connectivity information. None of this information is stored, but, instead, is generated as needed.

[0020] In one embodiment, m hierarchical mesh models ($m=8$, for example) for m discrete levels of detail (LODs) are used, with every LOD, l , up-sampling its adjacent coarser LOD, $l-1$, by two. To eliminate the need for storing these models (about 5 million vertices at the finest LOD), the vertices, edges and faces of these

- 4 -

LODs are carefully indexed such that their positions and connectivity at any LOD can be computed on demand.

[0021] In general, the closer the viewer is to a given patch, the more detail the viewer expects to see and vice versa. There is a correlation between the range of elevation and each of the predetermined mesh LODs and, based on this correlation, the most appropriate LOD mesh topology is selected. Basically, the farther the user is away, the coarser the mesh that is selected. The selection of the LOD is determined by the elevation h of the view point. Each LOD is associated with a minimum elevation threshold, $H_{\min}^l$. An appropriate LOD, l^* , is chosen such that

$$H_{\min}^{l^*-1} > h \geq H_{\min}^{l^*}.$$

[0022] A tessellation is done to determine the number of sample points inside each patch and to triangulate the patch with smaller triangles to achieve a continuous level of detail. Before tessellation, patches that are certain to be invisible to the viewer are removed to reduce computation cost. Invisible patch removal consists of two steps: back face culling and view frustum culling.

[0023] Hierarchical back face culling 16 is shown in Figure 2. The back face culling removes faces on the back side of the spherical mesh facing away from the viewer, depending on how far away the viewer is. Basically, the closer the viewer is to the patch, the less front faces the viewer sees. The unseen faces are culled to simplify the calculation. Since each more detailed mesh topology up-samples its predecessor by two, by adding one more vertex between every two vertices in each of longitude and latitude directions, only the vertices added in the next finer LOD that lie on the boundaries of the coarser LOD needs to be checked for back face culling, in one embodiment. Therefore, instead of performing back face culling directly on the chosen LOD with quadratic algorithm complexity $O(n^2)$, back face culling is done hierarchically by testing the visibility of all vertices at the coarsest LOD and incrementally testing added boundary vertices at each adjacent finer LOD up to the chosen LOD. The hierarchical back face culling has linear algorithm complexity $O(n)$, which is much faster than the quadratic complexity.

- 5 -

[0024] Back face culling produces a front face mesh that needs to be updated only when the position of the view point changes. A check at the switch 17 determines whether the view point has changed, which would determine if hierarchical back face culling needs to be performed to update the front face mesh. View frustum culling is then applied to generate a minimal mesh containing only faces potentially in the view frustum. In the view frustum culling 18, the view frustum in 3D space is determined. The three dimensional view frustum indicates what the viewer can see based on the viewer's selected viewing position, elevation, and viewing angle. Patches which are not going to intersect the view may be removed to simplify the calculations.

[0025] Back face culling and view frustum culling request per-patch maximum and minimum elevation values for visibility estimation. The hierarchical back face culling 16 uses the maximum elevation data from the fetch max HQT block 34. The view frustum culling 18 uses the minimum elevation from the block 32 and the maximum elevation from the block 34. The complete file of elevation data may be encoded into the maximum elevation, the minimum elevation, and the average elevation data. In some embodiments, the elevation data may be obtained from public sources, such as NASA's Global Digital Elevation Map (GDEM).

[0026] Once the surfaces have been reduced, the next step is to do the actual tessellation. This involves the invoking of the tessellation metric 22, which basically is the scheme for determining how many sample points are needed for each patch. Samples are added along the edges E outside the selected patch, as well as inside the patch P, as shown in Figure 3. The number of these samples or sample points S is determined by the tessellation metric 22. Then, once the samples have been chosen, the sample points are connected into triangles T that are ultimately sent to the renderer.

[0027] Once a mesh topology has been selected, more details can be added in patches of interest within that topology. Each patch is a quadrilateral of four vertices and four edges. The model may be simplified to use a planar patch in one embodiment.

- 6 -

[0028] The patch tessellation 20 adds the samples to the edges and inside the patch and connects the points into triangles that are sent to the renderer. The patch evaluation 24 computes the position and the normal vector of each of the sample points using displacement mapping.

[0029] Displacement mapping is the adjustment from a coarse mesh LOD to a fine detailed mesh LOD. In one embodiment, displacement mapping is the adjustment from the perfect sphere to the actual topography of elevation in the real world. Thus, as shown in Figure 4, each point **P** on the sphere of radius **R** is adjusting by a distance **D**, along a normal **n** of the sphere, to the actual position **S**, on the surface of the Earth. The new surface position and normal vector **n_S** may be determined using displacement mapping techniques.

[0030] Thus, returning to the tessellation metric 22, it decides the tessellation resolution of a patch, i.e. how many sample points **S** are needed in a patch, to achieve a desirable continuous level of detail. The tessellation resolution of a patch is determined by the tessellation resolutions along the edges of a patch and the tessellation resolution inside a patch.

[0031] The tessellation resolution along an edge of a patch determines the number of sample points needed on an edge of a patch. It is done by computing how many segments an edge should be divided into, based on a number of factors.

[0032] The first factor is the viewing distance to the edge in question. The points **P_{c1}** and **P_{c2}** are the end points of an edge of a patch displaced by the average elevation of the patch and transformed into the camera coordinate system. Since these end points are determined in the camera coordinate system, they correspond to the viewing distance from the view point to the edge. The tessellation metric determines the average of the viewing distances as the first factor.

[0033] The second factor is the viewing angle. The edge size on the screen is determined by the viewing angle, so more points may be needed for edges that appear along long edges based on viewing angle. The end points are normalized and then a dot product is taken to give the cosine value of the viewing angle. When

- 7 -

the viewing angle is small, the cosine value is large and if the cosine value is in the denominator of the edge tessellation equation $T(e)$, the number of samples will be smaller.

[0034] The next factor is the flatness error ε_{flat} . If the edge is less flat, then more points will be needed to accurately depict the edge. The flatness error factor takes into account the error from approximating the sphere using the selected LOD mesh and the error due to the amount of elevation change. In one embodiment, an empirical scaling factor λ may be used to combine the two error terms.

[0035] The next factor takes into account the X and Y direction sizes (F_x and F_y) of the view port. The bigger the view port, the more detail that is needed.

[0036] The next factor, called k , is a constant scaling factor to convert the results of the tessellation metric equation into the appropriate range.

[0037] Then the tessellation metric $T(e)$ determines how many segments each edge should be divided into:

$$T(e) = k \frac{\varepsilon_{flat}(f)(F_x + F_y)}{|\mathbf{P}_{c1} + \mathbf{P}_{c2}| \left(1 + \frac{|\mathbf{P}_{c1}|}{|\mathbf{P}_{c1}|} \bullet \frac{|\mathbf{P}_{c2}|}{|\mathbf{P}_{c2}|} \right)}$$

$\mathbf{P}_{c1}$ and $\mathbf{P}_{c2}$ are the 3D positions of an edge's two end points displaced by the average elevation of the patch and transformed into the camera's coordinate system.

$\frac{1}{2}|\mathbf{P}_{c1} + \mathbf{P}_{c2}|$ gives the average viewing distance of an edge. $\left(1 + \frac{|\mathbf{P}_{c1}|}{|\mathbf{P}_{c1}|} \bullet \frac{|\mathbf{P}_{c2}|}{|\mathbf{P}_{c2}|} \right)$ is the viewing angle computed from the dot product of the normalized two end points.

[0038] The flatness error of a patch, $\varepsilon_{flat}(f)$, i.e. the error of approximating a high order spherical patch by a flat face, is estimated by a weighted sum of the spherical approximation error of the chosen LOD, ε_{sphere}^l , and the topographic unevenness bounded by the maximum elevation difference between the maximum elevation $h_{max}(f)$ and minimum elevation $h_{min}(f)$ of the patch, which are derived from the GDEM elevation data in one embodiment,

- 8 -

$$\mathcal{E}_{flat}(f) = \mathcal{E}_{sphere}^l + \lambda[h_{max}(f) - h_{min}(f)].$$

[0039] The spherical approximation error of a LOD, $\mathcal{E}_{sphere}^l$, is computed, in advance at the time when the LOD is chosen, based on the maximum distance error between a flat face of the LOD mesh and its corresponding perfect spherical patch,

$$\mathcal{E}_{sphere}^l = \sqrt{R(1 - \sqrt{\cos^2 \frac{\theta^l}{2} - \sin^2 \frac{\theta^l}{2}})},$$

where R is the radius of the planet, and θ^l is the angle in the longitude or latitude direction that a single face of the LOD covers.

[0040] Once the tessellation resolution for each edge of a patch $T(e)$ is determined, the tessellation resolution for the inside (i.e. non-edge) points of a patch $T(f)$ by averaging the tessellation resolutions of all the edges of a patch,

$$T(f) = \text{avg}_{e \in f} T(e).$$

[0041] The tessellation resolution on each outside edge may be readjusted so that edges between adjacent patches are consistent to prevent visual artifacts such as edge cracking. Thus, the common edge determination for two adjacent patches is averaged to get the actual tessellation resolution $T(e)$ for the common edge,

$$T(e) = \text{avg}_{e \in f} T(f).$$

[0042] Vertices on patch corners may be shared by four or more visible patches. The tessellation resolution of each vertex $T(v)$ is determined by the average tessellation resolution of all the patches that share this vertex,

$$T(v) = \text{avg}_{v \in f} T(f).$$

The vertex tessellation resolution is not used for tessellation, but is used to determine the resolution of the topography data to use.

[0043] All the above calculations of tessellation resolutions $T(f)$, $T(e)$ and $T(v)$ are done in floating point values, in one embodiment. The final result of $T(f)$, $T(e)$ or $T(v)$, which may end up being an integer plus a fraction, is rounded up to the next higher integer that is a power of 2.

- 9 -

[0044] Thus, in one embodiment, the tessellation metric uses the maximum elevation data, the minimum elevation data, and the average elevation data, as indicated in Figure 2. Note that the corner points P_{c1} and P_{c2} are not the original corner points of the flat patch, but are actually the adjusted points for elevation (adjusted by average elevation of the patch). Therefore, the elevation data is needed in order to make the adjustment.

[0045] Given a patch with computed tessellation resolutions, a tessellator generates sub-sampling points with uv coordinates and triangular connectivity. The maximum tessellation resolution is 512, in one embodiment, equivalent to nine additional LODs.

[0046] Finally, the positions and normal vectors of the tessellated points are computed by displacement mapping elevation data during patch evaluation and fed to the rendering engine.

[0047] At each newly determined sample point, its three dimensional position and surface normal vector are computed using displacement mapping of the elevation data. The tessellation patch evaluation 24 for each patch produces a list of three dimensional coordinates of positions, a list of three dimensional normal vectors, and a list of indices representing the triangular connectivity of the sample points of the patch (block 26). The list of positions, normal vectors, and indices are then fed to OpenGL back end renderer 28 for rendering, in one embodiment.

[0048] The hierarchical back face culling, view frustum culling, tessellation metric, and patch evaluation all use elevation data. That data may be retrieved in real time by a server-client asynchronous data fetching system.

[0049] When elevation data is needed during hierarchical back face culling, view frustum culling, tessellation metric computation, and patch evaluation, the geometry processing pipeline (Figure 2) checks the memory cache 38 of the data client 40. If the needed data is not available, the geometry processing pipeline returns with the best data available in the cache 38 and resumes processing. Meanwhile, the data

- 10 -

client 40 sends a query over the network 50 to the data server 52 for the required data.

[0050] Once a query is received, the server 52 loads the corresponding type of elevation data 54, 56, 58 or 60 specified in the query from hard disk 92(Figure 8) and sends the data over the network 50 to the client 40, while storing a copy of the data in its own memory cache 62. Once the client receives the data from the server, the client stores the data in client memory cache 38 so that the data is available for geometry processing thereafter.

[0051] Hierarchical back face culling, view frustum culling, and tessellation metric computation require per-patch maximum, minimum and average elevation values, while patch evaluation requires detailed elevation data for displacement mapping. An efficient compression/decompression scheme is needed for real time retrieval of this large-scale data.

[0052] A hierarchical lossless compression algorithm may be applied to the elevation data offline to generate different types of elevation data needed by different parts of the geometry processing pipeline. The resulting data are stored in a hierarchical quadtree data structure. A hierarchical lossless decompression at run time allows for quick access to different types of elevation data in any region of interest at any resolution. The hierarchical data structure may be efficient for hard disk input/output on the server, for network transmission between the server and client, and for memory caching on both the server and the client.

[0053] Haar wavelets are fast to compute, but do not achieve a high compression ratio. A combination of Haar transform and bit packing is used for random data access and compression in one embodiment. Efficient data structures, file format, and a per-thread cache hierarchy are designed to provide a fast, scalable data encoding and retrieval solution despite the low compression ratio.

[0054] The basic 2D Haar transform involves the convolution of four 2x2 templates with the original data, which can be recovered by the same transform plus a scaling factor.

- 11 -

$$\begin{cases} X_{ij} = \begin{bmatrix} x_{00} & x_{01} \\ x_{10} & x_{11} \end{bmatrix} \otimes H_{ij} \\ x_{ij} = \frac{1}{4} \begin{bmatrix} X_{00} & X_{01} \\ X_{10} & X_{11} \end{bmatrix} \otimes H_{ij}, \end{cases}$$

where

$$H_{00} = \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}, H_{01} = \begin{bmatrix} 1 & -1 \\ 1 & -1 \end{bmatrix}, H_{10} = \begin{bmatrix} 1 & 1 \\ -1 & -1 \end{bmatrix}, H_{11} = \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}.$$

The Haar transform is essentially an averaging and differencing operation. When applied recursively, using the hierarchical quadtree layout, as shown in Figure 7, Haar transform concentrates low frequency components in the data, making it particularly suitable for a light-weight bit packing compression method.

[0055] Haar transform can be used directly to encode hierarchical per-patch average values. Unfortunately, Haar transform cannot be used to encode hierarchical per-patch maximum or minimum values. For that purpose, a maximum and minimum transform, extending the structure of Haar transform, may be used in some embodiments. A 2D maximum transform can be performed on a 2x2 data grid

$$\begin{bmatrix} x_{00} & x_{01} \\ x_{10} & x_{11} \end{bmatrix} \text{ as follows to yield a new data grid } \begin{bmatrix} X_{00} & X_{01} \\ X_{10} & X_{11} \end{bmatrix}:$$

1. Find the maximum value $x'_{00} = \max\{x_{00}, x_{01}, x_{10}, x_{11}\}$.
2. Exchange positions between x'_{00} and x_{00} to ensure that x'_{00} is at the top left corner of the new grid $\begin{bmatrix} x'_{00} & x'_{01} \\ x'_{10} & x'_{11} \end{bmatrix}$. For example, if $x'_{00} = x_{01}$, exchange positions

between x_{01} and x_{00} so that $\begin{bmatrix} x'_{00} & x'_{01} \\ x'_{10} & x'_{11} \end{bmatrix} = \begin{bmatrix} x_{01} & x_{00} \\ x_{10} & x_{11} \end{bmatrix}$.

3. Difference from the maximum value $\begin{bmatrix} X_{00} & X_{01} \\ X_{10} & X_{11} \end{bmatrix} =$

$$\begin{bmatrix} x'_{00}, & x'_{00} - x'_{01} \\ x'_{00} - x'_{10} & x'_{00} - x'_{11} \end{bmatrix}.$$

- 12 -

4. Concatenate the original position bits of x'_{00} to X_{11} . For example, if $x'_{00} = x_{01}$, concatenate the binary number "01" to X_{11} .

[0056] The inverse maximum transform is straightforward:

1. Extract the original position binary number "ij" of x'_{00} from X_{11} .

2. Restore values
$$\begin{bmatrix} x'_{00} & x'_{01} \\ x'_{10} & x'_{11} \end{bmatrix} = \begin{bmatrix} X_{00} & X_{00} + X_{01} \\ X_{00} + X_{10} & X_{00} + X_{11} \end{bmatrix}.$$

3. Restore positions by exchanging x'_{00} and x'_{ij} .

[0057] The maximum transform is essentially a maximizing and differencing operation. When applied recursively using the same hierarchy in Figure 7, the maximum transform concentrates low frequency components in the data and is suitable for use with bit packing compression method.

[0058] This maximum transform can be easily revised into a minimum transform by slightly modifying the maximizing and differencing operations, with the same benefit from the unified hierarchical framework. A 2D minimum transform can be

performed on a 2x2 data grid $\begin{bmatrix} X_{00} & X_{01} \\ X_{10} & X_{11} \end{bmatrix}$ as follows to yield a new data grid

$$\begin{bmatrix} X_{00} & X_{01} \\ X_{10} & X_{11} \end{bmatrix}:$$

1. Find the minimum value $x'_{00} = \min\{x_{00}, x_{01}, x_{10}, x_{11}\}$.

2. Exchange positions between x'_{00} and x_{00} to ensure that x'_{00} is at the top left corner of the new grid $\begin{bmatrix} x'_{00} & x'_{01} \\ x'_{10} & x'_{11} \end{bmatrix}$. For example, if $x'_{00} = x_{01}$, exchange positions

between x_{01} and x_{00} so that
$$\begin{bmatrix} x'_{00} & x'_{01} \\ x'_{10} & x'_{11} \end{bmatrix} = \begin{bmatrix} x_{01} & x_{00} \\ x_{10} & x_{11} \end{bmatrix}.$$

3. Difference from the minimum value
$$\begin{bmatrix} X_{00} & X_{01} \\ X_{10} & X_{11} \end{bmatrix} =$$

- 13 -

$$\begin{bmatrix} x'_{00} & x'_{01} - x'_{00} \\ x'_{10} - x'_{00} & x'_{11} - x'_{00} \end{bmatrix}.$$

4. Concatenate the original position bits of x'_{00} to X_{11} . For example, if $x'_{00} = x_{01}$, concatenate the binary number "01" to X_{11} .

[0059] The inverse minimum transform is straightforward:

1. Extract the original position binary number "ij" of x'_{00} from X_{11} .
2. Restore values $\begin{bmatrix} x'_{00} & x'_{01} \\ x'_{10} & x'_{11} \end{bmatrix} = \begin{bmatrix} X_{00} & X_{00} + X_{01} \\ X_{00} + X_{10} & X_{00} + X_{11} \end{bmatrix}$.
3. Restore positions by exchanging x'_{00} and x'_{ij} .

[0060] A quadtree unit is the result of Haar, maximum or minimum transform of 4 levels, which encodes 4 LODs of a 16x16 2D data grid. Assuming 32 bits per data entry to cover the range of elevation values and with compression, the size of a quadtree unit is no more than 1KByte. This is close to the optimum size for network transmission, as the typical size of a network packet is a little over 1KByte.

[0061] A multi-layer hierarchy of quadtrees is constructed bottom-up, as illustrated in Figure 6. All the quadtrees at the bottom layer are encoded by recursive Haar, maximum or minimum transform independently. Every 16x16 quadtrees are grouped together and their root node values are used as input to encode their parent quadtree. Note that the root node value of a child quadtree is duplicated in their parent quadtree and this gives two benefits. First, no more than one quadtree needs to be decoded when retrieving a single data point. Second, a quadtree with uniform data does not need to be stored, as the values can be retrieved from its parent quadtree. Each quadtree contains information indicating if each of its child quadtrees exists and, if so, where its data are stored on the hard disk. This avoids unnecessary data fetching and minimize hard disk seek at run time and speeds up hard disk I/O.

[0062] The multi-layer quadtree hierarchy is scalable to any number of LODs. To encode per-patch average, maximum and minimum values, the grid of the quadtree

- 14 -

hierarchy is aligned with the discrete LOD mesh models, as shown in Figure 5. In one embodiment, two layers of quadtrees cover the 8 discrete LODs of per-patch average, maximum and minimum data. The number of layers can be increased to match any given number of discrete LODs.

[0063] To encode the detailed topographic data for displacement mapping, the quadtree hierarchy starts from two quadtrees at the top layer corresponding to two hemispheres and spawn child quadtrees recursively until the highest resolution of the original data is covered. In one embodiment, 5 layers of quadtrees are used, equivalent to 20 LODs, to cover the ASTER Global Digital Elevation Map (GDEM) data from NASA, which has a 30m resolution in spatial location and a 1m resolution in elevation. Note that the grid of this quadtree hierarchy is not aligned with the discrete LOD mesh models in one embodiment. This is designed to take advantage of the quadtree hierarchy's inherent scalability to data of heterogeneous resolutions, as any subset of quadtrees can be easily replaced by quadtrees of different data with a different resolution. This is particularly useful since the topographic data for a planet such as the Earth usually come in different resolutions for different regions.

[0064] Once the quadtree hierarchy is constructed, every 3 layers (or the maximum number of layers, in the case of per-patch value encoding) of quadtrees that descend from a common ancestor quadtree is clustered. The statistics are gathered from all the quadtrees in the cluster to compute a bit-length table, which records the maximum number of bits required for each of 16x16 data entries of all the quadtrees in the cluster. All the quadtrees in the cluster are then compressed by bit-packing according to the resulting bit-length table and stored in a file. The bit-length table is stored in the root quadtree of the cluster for decompression during data retrieval.

[0065] The offline construction pipeline of the quadtree hierarchy is shown in Figure 5. The raw GDEM files are passed to Haar transform 70a. The raw GDEM files are also passed to a minimum transform 70b and to a maximum transform 70c. Then, in each case, bit packing is done at block 72 to produce the appropriate files 76. For encoding per-patch average, minimum and maximum values, the data

- 15 -

loaded from the raw GDEM files are arranged to align with the discrete levels of detail mesh models to produce quadtree hierarchies aligned with the discrete levels of detail. For encoding detailed topographic data, the data loaded from the raw GDEM files may or may not align with the discrete levels of detail mesh models.

[0066] The choice of 3 layers is empirically determined as a tradeoff between maximizing the bit-packing compression ratio and minimizing the storage of bit-length tables. A better approach to improve the compression ratio without significantly increasing the storage of bit-length tables would be to perform an adaptive clustering on an entire quadtree hierarchy to extract the most representative bit-length patterns for bit-packing.

[0067] In order to guarantee a real-time performance of planetary rendering, a server-client partitioned system is used for asynchronous data fetching at run time, and its pipeline, shown in Figure 8.

[0068] When the client 40 is launched, an initial fetch (block 13 in Figure 2) sends queries to the server 52 to load the top layers of the maximum, minimum, average and data quadtree hierarchies, encoded as described previously.

[0069] On the server 52 side, a group of server cache management threads 93 listen for queries 88 from clients. Whenever a query is received, a server cache management thread checks (diamond 80) the server cache 62 first. If the requested quadtree is already in the cache, it is assembled (block 86) into a network data packet 90 together with its corresponding bit-length table and sent to the client 40 via the network 50. Otherwise, the server cache management thread locates the quadtree based on information available from its cached parent quadtrees and loads it (block 82) from hard disk 92. The quadtree is then assembled (block 86) and sent to the client 40. A copy of the encoded quadtree is stored (block 84) in the server cache 62 for fast retrieval next time the same data is queried.

[0070] On the processor of the rendering client 40, there are two groups of threads, geometry processing threads 94 and client cache management threads 151. Whenever a data packet 90 is received from the server, a global cache

- 16 -

management thread extracts the encoded quadtree from the packet and unpacks (block 150) its data bits based on the bit-length table coming with the quadtree. The unpacked quadtree is then stored (block 152) in a global cache 38b.

[0071] Whenever elevation data are needed during geometry processing, whether per-patch maximum, minimum, average values or detail data values, a geometry processing thread 94 determines the geo-location (i.e. longitude and latitude) and resolution 154 of the data and fetches corresponding raw data quadtrees from its local cache 38a. If they are available, tri-linear interpolation 156 is performed to compute the actual values needed for geometry processing. If a check at diamond 158 determines that a requested quadtree is not in the local cache, the geometry processing thread checks (diamond 160) the global cache 38b. If it is available, the inverse Haar, maximum or minimum transform (block 162) is applied to the unpacked quadtree from the global cache and store (block 164) the 4 LODs of quadtree raw data in the thread's local cache 38a. If the requested quadtree is not in the global cache, a query 166 is sent to the server 52. Meanwhile, the best available quadtree in the global cache with the highest resolution is retrieved (block 168), inverse-transformed (block 162), and copied (block 164) to the thread's local cache 38a, and further geometry processing continues with the resulting raw data.

[0072] The reason for keeping data encoded on the server and decoding in a distributed manner at different stages on the client is to keep storage as compact as possible while minimizing run time computation cost. As the global cache on the client is shared by multiple threads, locking is necessary to prevent a race situation, which is why a per-thread local cache is used to alleviate the lock contention problem. Together with data tiling, the per-thread local cache layer provides the extra benefits of reducing the access frequency to the system memory and improving hardware cache efficiency on the CPU. To take advantage of the quadtree hierarchy, as shown in Figure 6, whenever a quadtree is requested and cached, all its ancestor quadtrees are retrieved and cached as well. This speeds up the retrieval of the file information, the existence information and the bit-length table of nearby quadtrees, as there is a high chance of needing data in neighboring regions.

- 17 -

[0073] The server could be replaced by a cloud of server machines where the encoded elevation data are stored in a randomly distributed manner on each of the server machines to speed up the network response during data retrieval. A pre-fetching mechanism can be added to the system based on user interactions from mouse and keyboard inputs to further improve the data retrieval speed. The compression and decompression algorithms and data structures described above can be applied to any large 2D data sets, for example, the terrain texture data, for real time retrieval.

[0074] In some embodiments, the geometric processing described above is performed in a CPU. Lighting and shading may be done in a graphics processing unit (GPU) using conventional techniques in one embodiment.

[0075] The computer system 130, shown in Figure 9, may include a hard drive 134 and a removable medium 136, coupled by a bus 104 to a chipset core logic 110. A keyboard and mouse 120, or other conventional components, may be coupled to the chipset core logic via bus 108. The core logic may couple to the graphics processor 112, via a bus 105, and the main or host processor 100 in one embodiment. The host processor may have multi-threading, single instruction multiple data (SIMD) and double precision, wider advanced vector extensions (AVX) capabilities in some embodiments. In one embodiment, the host processor implements the geometrical processing described herein. In another embodiment, the graphics processor may be integrated with the host processor. The graphics processor 112 may also be coupled by a bus 106 to a frame buffer 114. The frame buffer 114 may be coupled by a bus 107 to a display screen 118. In one embodiment, a graphics processor 112 may be a multi-threaded, multi-core parallel processor using single instruction multiple data (SIMD) architecture. In one embodiment, the graphics processor may be used for shading and lighting.

[0076] In the case of a software implementation, the pertinent code 139 may be stored in any suitable semiconductor, magnetic, or optical memory, including the main memory 132 or any available memory within the graphics processor. Thus, in one embodiment, the code to perform the sequences of Figures 2 and 8 may be

- 18 -

stored in a non-transitory machine or computer readable medium, such as the memory 132 or the graphics processor 112, and may be executed by the processor 100 or the graphics processor 112 in one embodiment.

[0077] Figures 2 and 8 are flow charts. In some embodiments, the sequences depicted in these flow charts may be implemented in hardware, software, or firmware. In a software embodiment, a non-transitory computer readable medium, such as a semiconductor memory, a magnetic memory, or an optical memory may be used to store instructions and may be executed by a processor to implement the sequences shown in Figures 2 and 8.

[0078] The graphics processing techniques described herein may be implemented in various hardware architectures. For example, graphics functionality may be integrated within a chipset. Alternatively, a discrete graphics processor may be used. As still another embodiment, the graphics functions may be implemented by a general purpose processor, including a multicore processor.

[0079] References throughout this specification to “one embodiment” or “an embodiment” mean that a particular feature, structure, or characteristic described in connection with the embodiment is included in at least one implementation encompassed within the present invention. Thus, appearances of the phrase “one embodiment” or “in an embodiment” are not necessarily referring to the same embodiment. Furthermore, the particular features, structures, or characteristics may be instituted in other suitable forms other than the particular embodiment illustrated and all such forms may be encompassed within the claims of the present application.

[0080] While the present invention has been described with respect to a limited number of embodiments, those skilled in the art will appreciate numerous modifications and variations therefrom. It is intended that the appended claims cover all such modifications and variations as fall within the true spirit and scope of this present invention.

- 19 -

What is claimed is:

- 1 1. 1. A method comprising:
2 providing a plurality of selectable mesh topologies of different levels of
3 detail, each mesh topology made up of abutting faces; and
4 enabling continuous levels of detail tessellation of the faces of a
5 selected mesh topology.
- 1 2. The method of claim 1 including enabling random access to
2 geometrical data at different locations and levels of detail in a mesh topology.
- 1 3. The method of claim 2 including organizing the geometrical data in
2 quadtrees, each quadtree including less than a kilobyte of data.
- 1 4. The method of claim 3 including building a layered quadtree data
2 structure from said quadtrees from the bottom layer up, each successively higher
3 parent layer formed from the root nodes of a previous child layer.
- 1 5. The method of claim 4 including accessing the structure from the top
2 down.
- 1 6. The method of claim 5 including putting information in said quadtrees
2 about their children quadtrees.
- 1 7. The method of claim 1 including computing tessellation resolution by
2 combining viewing distance, viewing angle, flatness error, and view port size.
- 1 8. The method of claim 7 including computing flatness error using
2 spherical approximation error and topographic unevenness error.
- 1 9. The method of claim 1 including using adaptive clustering to extract
2 representative bit-length tables for bit packing compression.

- 20 -

1 10. The method of claim 9 including applying hierarchical Haar transform to
2 hierarchical maximum and minimum transforms and therefore using a unified
3 quadtree hierarchy data structure for data compression/ decompression.

1 11. A non-transitory computer readable medium storing instructions to
2 enable a processor to:
3 select one of a plurality of mesh topologies at different levels of detail;
4 and
5 randomly access geometrical data at different locations and levels of
6 detail in a mesh topology.

1 12. The medium of claim 11 further storing instructions to provide a
2 plurality of selectable mesh topologies of different levels of detail, each mesh
3 topology made up of abutting faces and to enable continuous levels of detail
4 tessellation of the faces of a selected mesh topology.

1 13. The medium of claim 12 further storing instructions to organize the
2 geometrical data in quadtrees, each quadtree including less than a kilobyte of data.

1 14. The medium of claim 13 further storing instructions to build a layered
2 quadtree data structure from said quadtrees from the bottom layer up, each
3 successively higher parent layer formed from the root nodes of a previous child
4 layer.

1 15. The medium of claim 14 further storing instructions to access the
2 structure from the top down.

1 16. The medium of claim 15 further storing instructions to put information in
2 said quadtrees about their children quadtrees.

- 21 -

1 17. The medium of claim 11 further storing instructions to compute
2 tessellation resolution by combining viewing distance, viewing angle, flatness error,
3 and view port size.

1 18. The medium of claim 17 further storing instructions to compute flatness
2 error using spherical approximation error and topographic unevenness error.

1 19. The medium of claim 11 further storing instructions to use adaptive
2 clustering to extract representative bit-length tables for bit packing compression.

1 20. The medium of claim 19 further storing instructions to apply
2 hierarchical Haar transform to hierarchical maximum and minimum transforms and
3 therefore using a unified quadtree hierarchy data structure for data
4 compression/decompression.

1 21. An apparatus comprising:
2 a processor to build a layered quadtree data structure from a bottom
3 layer of quadtrees up, each successively higher parent layer formed from the root
4 nodes of a previous child layer; and
5 a storage coupled to said processor.

1 22. The apparatus of claim 21, said processor to provide a plurality of
2 selectable mesh topologies of different levels of detail, each mesh topology made up
3 of abutting faces and to enable continuous levels of detail tessellation of the faces of
4 a selected mesh topology.

1 23. The apparatus of claim 22, said processor to enable random access to
2 geometrical data at different locations and levels of detail in a mesh topology.

1 24. The apparatus of claim 23, said processor to organize the geometrical
2 data in quadtrees including less than a kilobyte of data.

- 22 -

1 25. The apparatus of claim 21, said processor to access the structure from
2 the top down.

1 26. The apparatus of claim 26, said processor to load information in said
2 quadtrees about their children quadtrees.

1 27. The apparatus of claim 22, said processor to compute tessellation
2 resolution by combining viewing distance, viewing angle, flatness error, and view
3 port size.

1 28. The apparatus of claim 27, said processor to compute flatness error
2 using spherical approximation error and topographic unevenness error.

1 29. The apparatus of claim 22, said processor to use adaptive clustering to
2 extract representative bit-length tables for bit packing compression.

1 30. The apparatus of claim 29, said processor to apply hierarchical Haar
2 transform to hierarchical maximum and minimum transforms and therefore using a
3 unified quadtree hierarchy data structure for data compression/ decompression.

1 / 7




FIG. 1




FIG. 3




FIG. 2

3 / 7

FIG. 4





FIG. 6

4 / 7




FIG. 5




FIG. 7

6 / 7




FIG. 8




FIG. 9