

Terrain rendering
Due: Thursday, March 18, 12 noon

1 Summary

The final project involves rendering outdoor scenes. Your task will be to take an implementation of the ROAM mesh-simplification algorithm and build an interactive terrain engine. In addition to implementing the basic rendering engine, you will be responsible for implementing several special effects.

Since naïve rendering of a heightfield mesh requires excessive resources (a 1025×1025 grid has 2^{21} triangles) we will use a *continuous level-of-detail* algorithm to reduce the mesh to a reasonable size. Specifically, we will use the *split-only* version of the ROAM algorithm. This algorithm works by taking an initial coarse-grain triangulation of the heightfield (*e.g.*, two triangles) and then iteratively refining it until a predefined triangle budget is reached.

2 Heightfields

Heightfields are a special case of mesh data structure, in which only one number, the height, is stored per vertex. The other two coordinates are implicit in the grid position. If s_h is the horizontal scale, s_v the vertical scale, and \mathbf{H} a height field, then the 3D coordinate of the vertex in row i and column j is $\langle s_h j, s_v \mathbf{H}_{i,j}, s_h i \rangle$ (assuming that the upper-left corner of the heightfield has X and Z coordinates of 0). By convention, the top of the heightfield is north; thus, assuming a right-handed coordinate system, the positive X -axis points east, the positive Y axis points up, and the positive Z -axis points south. The heightfield is typically represented as a linear array of height samples, with the i, j element at index $iw + j$, where w is the width of the heightfield. Because of their regular structure, heightfields are trivial to triangulate; for example, Figure 1 gives two possible triangulations of a 5×5 grid. For this project, we will use the ROAM algorithm, which uses the triangulation shown on the left of Figure 1. Note that the direction of the split for a triangle can be determined by the sum of the row and column indices of the upper-left corner. If the sum is even, then the split runs from the NW corner to the SE corner, and if the sum is odd, the split runs from the SW corner to the NE corner.

3 Input format

A terrain data set is represented as a directory containing various files that define the scene to be rendered. These files include




Figure 1: Heightfield triangulations

- `map` — this file contains information about the terrain data set, such as scale, feature locations, and the direction of the sun.
- `hf.png` — this file contains the height-field data as a 16-bit grayscale image.
- `color.png` — this file contains the vertex color information (see Section 4.1).
- `shadow.png` — the precomputed shadow texture (see Section 4.3).
- `water.png` — this file contains a depth map that gives the water depth for each grid location (see Section 7.3).
- `detail.png` — detail texture (see Section 4.3).
- `trees` — location of trees (see Section 7.1).
- `geysers` — location of geysers (see Section 7.2).

The `LoadTerrain` function will be used to load the map data.

```
Map_t *LoadTerrain (const char *terrain, Viewer_t *view);
```

This function takes the name of the *directory* containing the terrain data set and returns a *map* object, which contains in-memory versions of the data. It also initializes the initial camera position and direction.

4 Rendering

The first part of this project is to implement a basic renderer on top of the ROAM CLOD implementation. Your fragment shader will need to use both the color texture and shadow texture to compute the fragment colors.

4.1 Lighting

For this part, we will add a single directional light (the sun) to the scene, which is specified in the map file. Adding lighting means that you will need to specify normals for your triangles as you

render them. The colors assigned to a given vertex are defined by the `color.png` file.

4.2 Fog

Fog adds realism to outdoor scenes. It can also provide a way to reduce the amount of rendering by allowing the far plane to be set closer to the view. The map file format includes a specification of the fog density and color.

4.3 Detail texture

To make the surface of the terrain look more realistic, your implementation should blend in a *detail texture*. Each map has a `detail.png` file that contains the texture.

You use the detail texture, which is essentially a noise texture, to modulate the surface color close to the camera (say out to 80 to 100 meters). The texture is allocated as an RGB image, you may want to copy it to an RGBA representation with a 40-50% alpha channel, which will mute its effect somewhat. You can also use alpha blending to get a smooth transition from textured polygons to untextured ones.

5 ROAM

The ROAM algorithm is organized around a dynamic representation of triangle meshes called *triangle binary trees*. Figure 2 gives an example of a tree and Figure 3 shows the corresponding levels of triangulation. In the split-only version of this algorithm, we compute a new tessellation of the heightfield each frame by starting with the two triangles that cover the whole heightfield and then refining the mesh. We assign triangles a priority based on the benefit of refining them (*e.g.*, error metrics). Each triangle in the mesh has three neighbors (except for those triangles on the border) as is shown in Figure 4.

As can be seen from these figures, constructing a binary triangle tree can be done as a recursive splitting procedure. The trick is that we only want to split a triangle if the resulting mesh provides a visibly more accurate approximation of the height field. Thus, we modify the recursive splitting procedure to split the triangle with the highest priority, where priorities are a measure of the visual effect of not splitting. We use a limit of the number of triangles in the mesh to control the amount of rendering work we do. Thus, the pseudocode for the tessellation phase is

```
initialize the mesh to top two triangles
while (size of mesh < limit) {
    split highest priority triangle
}
```

Splitting a triangle requires splitting the triangle's base neighbor (otherwise a T-junction results), but it may also require presplitting the neighbor, when it is at a higher-level in the binary triangle tree. Figure 5 shows this situation.




Figure 2: Binary triangle tree




Figure 3: Binary triangle tree levels




Figure 4: Triangle neighbors

5.1 Hints

You can adjust the priority of triangles to eliminate detail where it is not needed and to enhance detail where it is needed. For example, triangles that lie wholly outside the view frustum should have minimum priority, while the triangle containing the camera should have maximum priority. Since the view is mostly horizontal, you can approximate the view frustum as a triangle in the XZ plane. Given the vehicle's heading (a vector in the XZ plane) and the horizontal field of view, one computes the line equations for the sides of the frustum and then uses these equations to assign low priorities to triangles outside the view. Figure 6 illustrates this optimization for a small mesh.

6 Camera controls

Your implementation should support controls for the view as described in the following table:

UP ARROW	accelerate
DOWN ARROW	brake
LEFT ARROW	turn left
RIGHT ARROW	turn right
f	toggle fog
l	toggle lighting
w	toggle wireframe mode
+	increase level of detail (by $\sqrt{2}$)
-	decrease level of detail (by $\sqrt{2}$)
q	quit the viewer

For the navigation controls, you will need to sample the state of the arrow keys (instead of just reacting to keyboard events). GLUT provides a function for registering a callback that gets called




Figure 5: Forcing splits

when a special key is released:

```
void glutSpecialUpFunc (void (*func)(unsigned int key, int x, int y));
```

Thus you will need two callbacks to keep track of the state of the arrow keys. Since holding down a key generates a repeated sequence of key events, which take time to service, you can disable key repeats using the following GLUT call:

```
glutIgnoreKeyRepeat (1);
```

7 Visual effects

As part of your project, you should implement visual effects to make the terrain more interesting. We list a number of possible visual effects below, but you should feel free to be creative. You should implement at least two of the effects.

7.1 Trees

A barren landscape is uninteresting, so to make it more attractive we can add some vegetation. A map can contain a `tree` file, which is a text file containing the locations of trees in the terrain. Use the function

```
Tree_t *LoadTrees (Map_t *map);
```

to load the array of tree information.

For the purpose of this project, we use a simple representation of trees based on four concentric tapered cylinders (see Figure 7). We can use `gluCylinder` to render each of the tree components; to keep the polygon count down, you should use multiple levels of detail in your rendering. You




Figure 6: View-frustum culling




Figure 7: Tree constructed from tapered cylinders

may also want to consider various culling techniques.

7.2 Geysers

Driving around a static wasteland is boring, so we populate the terrain with geysers. Information about the position and size of geyser's is stored in the `geyser` file. Use the function

```
Geyser_t *LoadGeysers (Map_t *map);
```

to load the array of geyser information.

The behavior of a geyser is controlled by three parameters: the frequency f , the duration d , and the average height h . A geyser erupts on average once every $\frac{1}{f}$ seconds with an average duration of d seconds. The average height of an eruption is s .

We model geysers using particle systems. The initial velocity vector of the particles should

have a cone-shaped distribution. You will want to set the average velocity to get the particles to the average height. You should also taper-off the system (both in particle velocity and number of particles) to get a gradual ending of the eruption. For example, let d be the duration of the geyser's eruption, then we can define the “size” of the geyser as a function of time t (assuming $t = 0$ is the start of the eruption):

$$s(t) = \begin{cases} he^{-(\frac{2t}{d})^2} & 0 \leq t \leq d \\ 0 & t > d \end{cases}$$

For example, $s(d) \approx 0.02h$. You can use this size function to influence the initial particle velocity, number of particles, and particle lifetimes. Feel free to play with your parameters and other decay functions.

7.3 Water animation

The surface of the lakes is flat and static. We can make it more interesting by adding waves to the water surface. The `water.pgm` file contains a depth map that gives the water depth for each grid location (a value of zero means dry land). Use the function

```
Elev_t *LoadWaterMap (Map_t *map);
```

to load the water map. This map can be used to construct meshes that represent the lake surfaces, which can then be animated.

7.4 Grass

In addition to using detail textures to make nearby terrain look more realistic, it is possible to add grass and small bushes to the scene. There are a number of shader-based techniques for rendering realistic looking grass.