

Compression and rendering of high resolution planetary scale digital elevation models

Enrico Gobbetti and Fabio Marton
CRS4, Pula, Italy – www.crs4.it/vic/
{Enrico.Gobbetti|Fabio Marton}@crs4.it

ABSTRACT

In this contribution, we illustrate a technique for incorporating aggressive compression methods in the BDAM adaptive resolution framework for terrain rendering. The new structure provides a number of benefits: simplicity of data structures, overall geometric continuity, efficient compression and fast construction times, real-time decompression and rendering with configurable variable level-of-detail extraction, and runtime detail synthesis.

1. Background and Motivation

Real-time 3D exploration of digital elevation models (DEM), which are computerized representations of a planet's relief, is one of the most important components in a number of practical applications. Nowadays, high accuracy DEMs contain billions of samples, thus vastly exceeding memory and graphics capability of nowadays graphics platforms. To cope with this problem, there has been extensive research on output sensitive algorithms for terrain rendering. At present time, the most successful techniques for efficiently rendering very large datasets are based on two competing approaches: adaptive coarse grained refinement from out-of-core multiresolution data structures (e.g., BDAM [EG03], P-BDAM [IEEEViz 03], Chunked ROAM [IEEEViz04]) and in-core rendering from aggressively compressed pyramidal structures (e.g., Geometry Clipmaps [SIGGRAPH04]). The first set of methods are very efficient in approximating a planar or spherical terrain with the required accuracy and in incrementally communicating updates to the GPU as the viewer moves, but their multiresolution structure footprints typically require out-of-core data management. The second kind of approach, limited to planar domains, uses nested regular grids centered about the viewer. This kind of approach ignores local adaptivity, but is able to exploit structure regularity to compress data so that it typically succeeds in fitting all data structures entirely in core memory, thereby avoiding the complexity of out-of-core memory management. In this contribution, we illustrate a technique for incorporating aggressive compression methods in the BDAM framework. The new structure provides a number of benefits: simplicity of data structures, overall geometric continuity, efficient compression and fast construction times, real-time decompression and rendering with configurable variable level-of-detail extraction, and runtime detail synthesis. The efficiency of the approach is demonstrated on a



number of test cases, including the compression and interactive visualization of the whole planet Mars created from high resolution Mars Orbiter Laser Altimeter data.

2. Technique Summary

Terrain LOD algorithms use a hierarchy of mesh refinement operations to adapt the surface tessellation. The main idea behind the Batched Dynamic Adaptive Meshes (BDAM) approach is adopt a more complex primitive than the single vertex or triangle to define the structure: a small surface patches composed of a batch of a few thousands of triangles. The benefits of this approach are that the per-triangle workload to extract a multiresolution model is highly reduced and the small patches can be preprocessed and optimized off-line for a more efficient rendering. We summarize here the main concepts behind BDAM. Please refer to the original papers [EG03, IEEEVIZ04] for further details..

In BDAM, the small patches form a hierarchy of right triangles (HRT) that is coded as a binary tree. This representation can be used to easily extract a consistent set of contiguous triangles which cover a particular region with given error thresholds. These small triangular patches can be batched (hence the name) to the graphics hardware in the most efficient way. Therefore, each bintree node contains a small chunk of contiguous well packed tri-stripped triangles. To ensure the correct matching between triangular patches, BDAM exploits the right triangle hierarchy property that each triangle can correctly connect to: triangles of its same level; triangles of the next coarser level through the longest edge; and triangles of the next finer level through the two shortest edges.

To guarantee the correct connectivity along boundaries of different simplification levels, triangular patches are built by constraining simplification to keep the vertices on the shortest edges fixed. This way, each mesh composed by a collection of small patches arranged as a correct bintree triangulation still generates a globally correct triangulation. In practice, the consistency condition requires for each patch that vertices along the longest edge are shared with patches of the same level or with patches at the next coarser level.

In the original BDAM work, each patch was represented as a TIN, with inner vertices free to be located at arbitrary positions inside the patches. In order to efficiently compress the data, we introduce the following modifications in the framework: (a) all vertices are arranged on a regular grid, and a $\sqrt{3}$ pattern is used to move from a level to the next; (b) all patches share the same connectivity pattern, which is optimized as a cache coherent triangle strip; (c) the data stored for a given patch consists in the difference from the prediction obtained by applying an interpolatory subdivision step from the coarser level and the actual data; (d) difference data is wavelet coded and quantized; (e) at run-time, the required patches are reconstructed on-the-fly by a refinement method that keeps in-core all reconstructed data above the current cut.

Tests on large scale datasets, such as the full reconstruction of planet Mars from high resolution MOLA Laser Altimeter data, illustrate that this technique is able to provide compression rates similar to those of Geometry Clipmaps (100x compression with 10% of grid spacing RMS error), while being more general and fully adaptive.

3. Conclusions

We have briefly discussed our current work in the domain of compression and rendering of high resolution planetary scale digital elevation models. The talk will illustrate our results with live demos. Technical papers on the subject are in preparation.

4. Acknowledgments

This research is partially supported by the CRIMSON project (EU-FP6-510166), funded under the European FP6/IST program. The source 3D datasets are provided by and used with permission of the NASA MOLA team. Work on batched multiresolution triangulations is carried out in cooperation with ISTI-CNR.

References

More information on this research is available on the web at the following URL

- <http://www.crs4.it/vic/> - CRS4 Visual Computing Group