Geophysical Research Abstracts,
Vol. 11, EGU2009-3951-4, 2009
EGU General Assembly 2009

Ray Casting of Large Multi-Resolution Volume Datasets

C. Lux (1) and B. Fröhlich (2)
(1) (christopher.lux@medien.uni-weimar.de), (2) bernd.froehlich@medien.uni-weimar.de

High quality volume visualization through ray casting on graphics processing units (GPU) has become an important approach for many application domains. We present a GPU-based, multi-resolution ray casting technique for the interactive visualization of massive volume data sets commonly found in the oil and gas industry.

Large volume data sets are represented as a multi-resolution hierarchy based on an octree data structure. The original volume data is decomposed into small bricks of a fixed size acting as the leaf nodes of the octree. These nodes are the highest resolution of the volume. Coarser resolutions are represented through inner nodes of the hierarchy which are generated by down sampling eight neighboring nodes on a finer level. Due to limited memory resources of current desktop workstations and graphics hardware only a limited working set of bricks can be locally maintained for a frame to be displayed. This working set is chosen to represent the whole volume at different local resolution levels depending on the current viewer position, transfer function and distinct areas of interest. During runtime the working set of bricks is maintained in CPU- and GPU memory and is adaptively updated by asynchronously fetching data from external sources like hard drives or a network. The CPU memory hereby acts as a secondary level cache for these sources from which the GPU representation is updated.

Our volume ray casting algorithm is based on a 3D texture-atlas in GPU memory. This texture-atlas contains the complete working set of bricks of the current multi-resolution representation of the volume. This enables the volume ray casting algorithm to access the whole working set of bricks through only a single 3D texture. For traversing rays through the volume, information about the locations and resolution levels of visited bricks are required for correct compositing computations. We encode this information into a small 3D index texture which represents the current octree subdivision on its finest level and spatially organizes the bricked data. This approach allows us to render a bricked multi-resolution volume data set utilizing only a single rendering pass with no loss of compositing precision. In contrast most state-of-the art volume rendering systems handle the bricked data as individual 3D textures, which are rendered one at a time while the results are composited into a lower precision frame buffer. Furthermore, our method enables us to integrate advanced volume rendering techniques like empty-space skipping, adaptive sampling and preintegrated transfer functions in a very straightforward manner with virtually no extra costs.

Our interactive volume ray tracing implementation allows high quality visualizations of massive volume data sets of tens of Gigabytes in size on standard desktop workstations.