Chapter 6

Summary and Concluding Remarks

SAMR (Structured Adaptive Mesh Refinement), introduced in 1989 by Berger and Collela [9], is a special adaptive method for solving partial differential equations. In this approach the computational domain is initially covered by a set of coarse, structured subgrids. During the computation local error estimators are utilized to detect cells that require higher resolution. These cells are covered by a set of axis-aligned, structured subgrids, which overlay the refined regions of the coarse base grid.

In this thesis we developed direct and indirect volume visualization algorithms for scalar fields that are defined on SAMR grids. In particular we investigated the applicability of SAMR data structures for volume rendering of large datasets. Nowadays state-of-the-art volume rendering approaches leverage the powerful texture units of modern graphics hardware, which perform fast bi-, respectively trilinear interpolation of texture samples and thereby allow interactive frame rates even on standard consumer PC systems, compare Section 3.5.

But due to the enormous rate of increase in resolution of datasets, which are for example generated by 3D imaging devices or numerical simulations, performance is still an issue even when utilizing highly specialized graphics hardware. The performance limiting factors are the fill rate of the hardware, i.e. the number of pixels values that can be computed per time unit, the available amount of texture memory, and the I/O-bandwidth. In Section 4.4 of Chapter 4 we investigated approaches that accelerate texture-based volume rendering for the frequently occurring case of large, but sparse data, i.e. highly resolved data where only a small fraction of the voxels contains relevant information. The relevance criterion might for example be given by voxel transparency or by some cardinal function, e. g. choosing material subsets of segmented data.

Our strategy was to restrict most of the rendering work to the relevant parts of the volume. In order to benefit from the strength of texture hardware, the resulting coverage has to consist of few axis-aligned, non-overlapping regions, that can be processed separately in a view-consistent order for each viewpoint. We presented an approach, in which these regions are covered by subvolumes that are leaf nodes of a SAMR tree representation of the data volume. The tree was constructed by hierarchical clustering of cells that contain relevant data.

The AMR approach obtains a good balance between the number of boxes needed for tightly capturing the relevant voxels of the data volume and the number of covered non-relevant ones. The amount of created texture bricks is much smaller than the number of bricks created by an alternative octree approach, resulting in significant performance gains for the AMR approach, especially for the sparse datasets.

In Section 4.3 we further discussed a raycasting approach that is based on the utilization of an adaptive integration scheme. In particular for smooth datasets this results in a considerable reduction of interpolation operations and in an increase of the achieved rendering performance compared to the standard approach.

In Chapter 5 we described our work in the field of visualizing SAMR data. We discussed the topics of fast point location and spatial interpolation of discrete SAMR data in Sections 5.3 and 5.2. In Section 5.4 we presented indirect volume rendering methods for artifact-free visualization of carpet plots and planar slices as well as an algorithm for extracting C^0-continuous isosurfaces from data defined on unrestricted SAMR grids, i. e. grids that contain adjacent cells that differ by more than one level of refinement.

We further presented a hardware accelerated volume rendering approach for SAMR data that utilizes 3D textures in Section 5.5. Since current texture hardware requires axis-aligned texture bricks which contain cells on the same resolution level, some preprocessing was necessary. For this we proposed a decomposition of the data domain into disjoint axis-aligned areas of constant cell size that results in a small number of blocks. Since the partition maintains the multi-resolution representations of the grid function in the refined areas, the traversal of sub-branches of the tree can be pruned in the rendering phase, based on a view-dependent projection criterion.

Generating smooth animations of time-dependent data requires dense output. If the frequency of available time steps is too low, dense output is usually obtained by some sort of temporal interpolation between the given key-frames. Creation of dense output from a given set of time steps via interpolation is especially problematic for AMR data, due to potential changes of the underlying grid topology during the time evolution and the fact that usually not all generated subgrids are stored by the simulation. These problems are related to the fact that in AMR approaches the subgrid structure generally changes after each update, due to a regridding procedure that aims to place grids in regions where higher resolution is required. The time step size is usually halved between two consecutive levels of refinement, i.e. it increases exponentially, so storing all intermediate subgrids would result in huge amounts of data. In Section 5.6 we proposed an approach for temporal interpolation of AMR data that addresses these problems. In order to handle the problem of varying grid topology during time evolution, intermediate grid hierarchies are generated by merging the grids on the corresponding refinement levels. In a second step a nested grid structure is induced, employing a clustering algorithm. We investigated different interpolation schemes like linear, Hermite or flux-based interpolation, that allows a better approximation of the underlying grid function in the case that the data represents a conserved quantity. The overall algorithm is fast and allows an on-the-fly generation of interpolated frames.

In summary, we developed algorithms for

- accelerated software-, and hardware-based volume rendering of 3D scalar data,
- planar slicing and the display of height fields for cell-, and vertex-centered data on unrestricted SAMR grids,
- C^0-continuous isosurface extraction for cell-, and vertex-centered data on unrestricted AMR grids,
- software-, and hardware-based direct volume rendering for cell-, and vertex-centered data on unrestricted SAMR grids and
- temporal interpolation of SAMR data with varying grid topology.

All algorithms have been implemented within the visualization framework AMIRA. They have been employed in several visualization projects and are now used for data analysis by several international research groups in the fields of cosmology and numerical relativity.