Horst Joachim Mayer
Dr. sc. hum.

Interactive volume visualization in client/server-environments

Geboren am 21.10.1968 in Kronberg/Ts.
Reifeprüfung am 25.5.1987 in Gießen
Studiengang der Fachrichtung Medizinische Informatik vom WS 1988 bis SS 1994
Vordiplom am 11.10.1990 an der Universität Heidelberg/ Fachhochschule Heibronn
Diplom am 20.6.1994 an der Universität Heidelberg/ Fachhochschule Heibronn

Promotionsfach: Deutsches Krebsforschungszentrum (DKFZ)
Doktorvater: PD Dr. H.-P. Meinzer

The goal of this work was to make volume visualization applicable for use in interactive environments. New methods have been developed in several fields of research: sequential optimizations, parallelization, accessibility over a network, and realtime hybrid visualization. This work is structured according to those topics.
Defining objects in medical data sets is usually a prerequisite for visualization and, conversely, visualization can help in the segmentation process by providing efficient feedback of intermediate results to the user. Therefore, the guideline for the development of new methods was to allow fast switching between image processing and visualization.
Direct volume visualization is very compute and memory-access intensive, due to the large amounts of data that have to be processed. A plethora of methods have been described in the literature to speed up the rendering process. The most important ones have been summarized in this work. However, most methods improve image generation speed by precalculating properties such as gradients, shading, mapping, or distance maps. New methods have been developed that avoid lengthy preprocessing steps to speed up the cycle of image processing and visualization in an integrated system.
A fast background traversal method has been developed that is based on two nested sampling loops. Only nearest neighbor interpolation is used in empty areas of the volume data. Switching between the loops is fast and, unlike in other methods of this type, requires no preprocessing to retrieve information about object boundaries. Depending on the contents of the volume data, an approximate 30% gain in performance is achieved as compared to ordinary raycasting.
Space leaping tries to avoid some of the background traversal altogether, by jumping over empty image areas. This work describes for the first time, a generalization of the original idea for C-buffer (coordinates buffer) space leaping that works for combinations of translations and rotations and for perspective raycasting. A correction algorithm is applied that detects holes and occlusions in a relief of start coordinates. It is shown that the coordinates buffer has to be divided into sectors, each one requiring a different scan direction for the traversal of the intermediate coordinate buffer during relief correction. The new method correctly and efficiently identifies areas where new structures can become visible that have previously been occluded. Both aspects have not been handled in any previous work in this field. A novel approach to the background problem is presented that is based on the generation of coordinates on the volume back faces using a scan conversion algorithm. Furthermore, a new method has been developed to reuse information from the surrounding to close gaps in the projected relief. An additional performance gain of nearly 50% is achieved by using C-buffer space leaping to generate successive images in an animations.

The description of the Heidelberg Raytracing Model (HRM) includes a shadow casting light source. The calculation of the shadows could be integrated into the raycasting process, introducing almost no overhead over simple raycasting.
The new volume visualization algorithm, integrating all of the described speedup techniques, achieved interactive frame rates (2 frames per second) on a single R10000 processor. The frame rates can compete with the most advanced volume visualization algorithms published. However, unlike in those algorithms, no preprocessing is required.
An integrated parallel image processing and visualization system on distributed memory parallel architectures has been developed. The design of the system trades efficiency of separate operators for best overall performance by using a common data partitioning scheme. Thus, redistribution of volume data is avoided, because it was found to negatively affect the performance of such a system even on high performance parallel computers. Benchmarks were performed on a distributed memory IBM SP/2. All algorithms did scale well up to four processors, the image processing operators even up to eight. The PVM version used was identified to be a major limiting factor for the scalability of the system. Performed tests indicated, that especially the scalability of the parallel volume visualization would be considerably improved by using PVMe, a version of PVM optimized for the IBM SP/2.
To prove the concept of the parallelization, a complex image analysis system has been reimplemented that requires the whole parallel operator set. Using eight processors, the runtime could be reduced from approximately 1.5 hours for the serial version to less than 25 minutes, required to be applicable in a clinical routine workflow. The parallel image processing and visualization functions have also been made accessible interactively from a segmentation and visualization tool integrated into an existing teleradiology system.
Powerful systems, either serial or parallel, with large memory are needed to host segmentation and visualization. To provide access to, and allow resource sharing of such computers, a network interface has been developed that is specifically tailored to the work with volume data. Unlike with the common RPC paradigm, the data is not transferred over the network for every processing request. Instead, a server generates and manages handles for the data. The server design also includes a memory management strategy, adapted to the requirements of interactive processing.
A new hybrid visualization approach has been developed. Volume visualized images can be merged with surface based graphics inside the OpenGL environment. The surface based graphics in the hybrid scene can be manipulated in realtime. This provides the basis for numerous multi-object interaction models, like pointers or cutplanes in the scene, which would not have been possible with volume visualization alone. Hybrid visualization can also offload the server and the network by allocating time critical graphical interactions on the client side. Realtime interactions can be used to set parameters for calculations on the server.
It can be concluded that some major breakthroughs could be achieved, allowing the setup of truly interactive image processing and visualization systems in a way not possible before. The new methods can work in concert to provide the features necessary to build efficient interactive systems. New techniques for fast volume visualization provide interactive frame rates, such that viewing and mapping parameters can be adjusted interactively. A new hybrid visualization approach allows additional graphical elements to be employed to interactively navigate in the scene, even if computing the volume visualization should be too slow to be manipulated directly. Scalable parallel processing methods have been developed for application areas that require fast image processing to meet given time constraints, thus enabling a successful integration of volume visualization and segmentation applications in routine use. Finally, a client/server architecture provides efficient access to high performance computers and allows to share computing resources.