

Recent results in rendering massive models on horizontal parallax-only light field displays

M. Agus¹, F. Bettio¹, F. Marton¹, G. Pintore¹, and A. Zorcolo¹

¹ CRS4, Pula, Italy, vic@crs4.it

Abstract—In this contribution, we report on specialized out-of-core multiresolution real-time rendering systems able to render massive surface and volume models on a special class of horizontal parallax-only light field displays. The displays are based on a specially arranged array of projectors emitting light beams onto a holographic screen, which then makes the necessary optical transformation to compose these beams into a continuous 3D view. The rendering methods employ state-of-the-art out-of-core multiresolution techniques able to correctly project geometries onto the display and to dynamically adapt model resolution by taking into account the particular spatial accuracy characteristics of the display. The programmability of latest generation graphics architectures is exploited to achieve interactive performance. As a result, multiple freely moving naked-eye viewers can inspect and manipulate virtual 3D objects that appear to them floating at fixed physical locations. The approach provides rapid visual understanding of complex multi-gigabyte surface models and volumetric data sets.

I. INTRODUCTION

MULTISCOPIC visualization is an emerging display technology that aims to reproduce three dimensional scenes by generating observer independent light fields. The accurate reproduction of 3D light fields requires generating a large number of light beams of appropriate origin, direction, and color. Recent advances in 3D display design have demonstrated that interactive high resolution light field display technology is practically achievable. Even though considerations on human vision specifics can drastically reduce the amount of data that has to be encoded in a reconstructed

light field, rendering still remains a complex and computationally intensive task, which has limited until very recently the applicability of such displays to presentation of static images, prerecorded movies, or small graphics models.

In this contribution, we report on specialized out-of-core multiresolution real-time rendering systems able to render massive surface and volume models on a special class of horizontal parallax-only light field displays. The rendering methods employ state-of-the-art out-of-core multiresolution techniques able to correctly project geometries onto the display and to dynamically adapt model resolution by taking into account the particular spatial accuracy characteristics of the display. As a result, multiple freely moving naked-eye viewers can inspect and manipulate virtual 3D objects that appear to them floating at fixed physical locations. The approach provides rapid visual understanding of complex multi-gigabyte surface models and volumetric data sets. The rest of the paper is organized as follows: section II briefly describes the design of the light field displays considered, while section III and section IV detail the rendering systems developed to interactively visualize, respectively, massive surface models and scalar volume datasets.

II. LIGHT FIELD DISPLAY CONCEPT

In the light field displays considered here [1], projectors are densely arranged in a horizontal linear array behind the screen, each one projecting a specific image onto the holographic

screen to build up a light field (see figure 1 left). Mirrors located at the sides of the display reflect back onto the screen the beams that would otherwise be lost, creating in this way virtual projectors that increase the display field of view. Each projector emits light beams toward a subset of the points of the holographic screen, so that each screen point is hit by multiple light beams coming from different projectors. The screen, a holographically recorded, randomized surface relief structure, performs selective directional transmission of light beams. Horizontally, the surface is sharply transmissive, in order to maintain a sub-degree separation between views. Vertically, the screen scatters widely, so that the projected image can be viewed from essentially any height. The angular light distribution profile introduced by the holographic screen is characterized by a wide plateau and steep Gaussian slopes precisely overlapping in a narrow region in the horizontal direction. This results in a homogeneous light distribution and continuous 3D view with no visible crosstalk within the field of depth determined by the angular resolution. For building a rendering pipeline on this kind of displays it is necessary to determine where 3D points should be drawn on a given projector to produce a perspective correct image for the viewer. However, the linear perspective is not sufficient, because it ignores the transformation performed by the holographic screen. Since the screen is selective only in the horizontal direction, but scatters widely in the vertical one, the displayed light field’s dimensionality is reduced, and the application must decide how to deal with the missing degree of freedom. As explained in [2], the solution consists of using a multiple center of projection (MCOP) approach (see figure 1 right).

III. MULTIREOLUTION VISUALIZATION OF MASSIVE MODELS

In this section, we briefly describe our method to drive the light field display in order to get scalable visualization of massive models. The technique is a parallel spatial 3D display-aware version of Adaptive TetraPuzzles [3].




Fig. 1. **Display concept.** Left: Each projector emits light beams toward a subset of the points of the holographic screen. Top right: horizontally, the screen is sharply transmissive and maintains separation between views. Bottom right: vertically, the screen scatters widely, so that the projected image can be viewed from essentially any height, and the MCOP approximation chooses the ray passing through a virtual observer position.

It uses a distributed image generation system implemented on a cluster, with a front-end client PC selecting the level of detail from the multiresolution structure and multicasting graphics commands to back-end PCs. The characteristics of multiresolution techniques based on coarse grained adaptation are exploited to efficiently distribute data to back-end nodes as well as to efficiently pass them to the GPU through preferential paths.

The overall architecture of the rendering system is depicted in figure 2.

The Tetrapuzzles method uses a conformal hierarchy of tetrahedra generated by recursive longest edge bisection to spatially partition the model in a preprocessing step. Each tetrahedral cell contains a precomputed simplified version of the original model. The representation is constructed off-line during a fine-to-coarse parallel out-of-core simplification of the surface contained in diamonds (sets of tetrahedral cells sharing their longest edge). Appropriate boundary constraints are introduced in the simplification process to ensure that all conforming selective subdivisions of the tetrahedron hierarchy lead to correctly matching surface patches (see [3] for more details). The main advantage of the method is its ability to rapidly produce seamless variable accuracy reconstructions by




Fig. 2. **Scalable rendering architecture.** A front-end client PC selects the level of detail from the multiresolution structure and multicasts graphics commands to back-end PCs that perform the actual rendering. Object-based communication with extensive caching leads to an efficient implementation.



Fig. 3. **Interaction sequence.** These images, taken from the accompanying video, show successive instants of interactive manipulation of the multi-million triangles colored datasets on the 50Mpixel display.

assembling precomputed patches.

Since rendering for the spatial 3D display requires the adaptation of vertices to very small (voxel-sized) triangles, controlling triangle shapes during simplification to reduce triangle counts in nearly flat areas is no longer important. Thus, instead of performing high-quality (quadric based) simplification [3], we construct diamonds with a simplification

method that produces (roughly) uniformly tessellated meshes, and use edge length as a measure of tessellation accuracy. This approach allows us to manage colored meshes by simply using a color-per-vertex representation. With such a structure, variable resolution rendering is implemented by simple stateless top-down traversals of the binary trees used to encode the tetrahedron hierarchy, which combine view-

frustum and contribution culling. The traversal is performed once per spatial 3D frame, and consequently generates as a result the set of patches that needs to be rendered for all the views. Representative video frames are shown in figure 3. The sequences were recorded with a hand held video camera, freely moving in the display’s workspace. Note the parallax effects and the good registration between displayed object space and physical space, which demonstrate the multi-user capability of the display. The perceived image is fully continuous. This is qualitatively very different from other contemporary multiview display technologies, which force users into approximately fixed positions, because of the abrupt view-image changes that appear at the crossing of discrete viewing zones [4].

IV. VISUALIZATION OF MASSIVE SCALAR VOLUMES

In this section, we briefly describe our method to drive the light field display in order to get scalable interactive volume ray casting visualization of huge volumetric datasets. Our technique is derived from a single pass out-of-core raycaster introduced in [5], adapted to work on light field displays similarly to [2]. The overall approach is based on a GPU raycaster, which follows rays generated by a MCOP projection model, while adaptively sampling pre-filtered versions of the dataset at resolutions matching the varying spatial accuracy of the display. In order to allow the volume rendering of very large datasets, we rely on an adaptive technique based on the decomposition of a volumetric data set into small cubical bricks, which are then organized into an octree structure maintained out-of-core. The octree contains the original data at the leaves, and a filtered representation of children at the inner nodes. Each node also stores the range of values, as well as high quality precomputed gradients. In order to efficiently support linear interpolation, we replicate one layer of neighboring samples at each brick boundary. The octree is stored in an out-of-core structure, based on Berkeley

DB, and data is losslessly compressed with the LZO compression library. The system has been developed in order to fulfill all requirements involved in the analysis of high quality and high resolution medical and anatomical data. Volume data are represented as 16 bit scalars, and 32 bit gradients. The gradients are precomputed by employing high quality 5x5x5 Sobel filtering. The system is also able to manage and render segmented datasets represented as labelled voxels. In that case, the precomputation of levels of detail is modified to choose for each value the most popular label instead of the average. At runtime, a working set of bricks is generated and incrementally maintained on CPU and GPU memory by asynchronously fetching data from the out-of-core octree. The working set is created by an adaptive loader on the basis of the current object position in relation to the screen and transfer function. Once the current working set is defined, a compact indexing structure, which spatially organizes the current working set into an octree hierarchy, is encoded in a small texture. This structure spatially organizes the leaves of the current view dependent representation into an octree with neighbor pointers. The inner nodes of this structure simply contain pointers to children, and only the leaves refer to volume data nodes stored in the memory pool. The spatial index structure is exploited by an efficient stackless GPU raycaster, which is able to compute the volume rendering integral by enumerating non-empty bricks in front to back order, adapting sampling density to brick resolution, and stopping as soon as the accumulated opacity exceeds a certain threshold, updating both the framebuffer and the depthbuffer. The prototype volume ray caster implements a number of composition strategies, including Direct Volume Rendering with a Phong illumination model, boundary enhancement and view dependent transparency [6]. When dealing with labeled data, trilinear filtering is substituted with nearest neighbor to preserve label values [7]. In order to drive the light field display the scene is rendered once per projector view. The resulting system is capable to interactively explore extremely large datasets on light field




Fig. 4. **Inspecting raw and segmented datasets on the light field display.** Top: Raw whole body contrast CTA acquired on a 16 detector CT scanner (source: Radiology Department, University of Geneva). The volume has a resolution of $512 \times 512 \times 1559$ with 16 bit/sample. Bottom: segmented leg reconstructed from MRI acquisitions (source: MiraLab, Geneva). The volume has a resolution of $404 \times 474 \times 2050$. Pictures are taken with a hand-held camera at different viewing angles, in order to highlight the horizontal parallax of the light field display. Images are rendered at about 6.5fps.

displays (see figure 4 for some examples). The prototype can be considered a testbed in a development process aimed at creating a really good 3D light field based radiology workstation in the future [8]. Perceptual evaluation tests have been carried out and tend to prove that the system provides correct depth cues, helps in layout discrimination and is clearly superior to two-dimensional displays for path tracing tasks, common in operations such as the understanding of vascular structures. The first feedback received from physicians and radiologists seems to confirm this fact. Future work aims at performing an evaluation of the system for diagnostic tasks in a clinical context.

REFERENCES

- [1] T. Balogh, T. Forgas, T. Agocs, O. Balet, E. Bouvier, F. Bettio, E. Gobbetti, and G. Zanetti, "A scalable hardware and software system for the holographic display of interactive graphics applications," in *EUROGRAPHICS 2005 Short Papers Proceedings*, 2005.
- [2] M. Agus, E. Gobbetti, J. Iglesias Guitián, F. Marton, and G. Pintore, "Gpu accelerated direct volume rendering on an interactive light field display," *Computer Graphics Forum*, vol. 27, no. 3, pp. 231–240, 2008, proc. Eurographics 2008.
- [3] P. Cignoni, F. Ganovelli, E. Gobbetti, F. Marton, F. Ponzio, and R. Scopigno, "Adaptive tetrapuzzles: Efficient out-of-core construction and visualization of gigantic multiresolution polygonal models," *ACM Trans. on Graphics (SIGGRAPH 2004)*, vol. 23, no. 3, pp. 796–803, 2004.
- [4] W. Matusik and H. Pfister, "3D TV: a scalable system for real-time acquisition, transmission, and autostereoscopic display of dynamic scenes," *ACM Transactions on Graphics*, vol. 23, no. 3, pp. 814–824, Aug. 2004.
- [5] E. Gobbetti, F. Marton, and J. Iglesias Guitián, "A single-pass GPU ray casting framework for interactive out-of-core rendering of massive volumetric datasets," *The Visual Computer*, vol. 24, no. 7-9, pp. 797–806, 2008, proc. CGI 2008.
- [6] S. Bruckner and M. E. Gröller, "Style transfer functions for illustrative volume rendering," *Computer Graphics Forum*, vol. 26, no. 3, pp. 715–724, Sept. 2007.
- [7] M. Agus, A. Giachetti, E. Gobbetti, J. Iglesias Guitián, and F. Marton, "Towards advanced volumetric display of the human musculoskeletal system," in *Eurographics Italian Chapter Conference*. Conference held in Salerno, Italy: Eurographics Association, June 2008.
- [8] M. Agus, F. Bettio, A. Giachetti, E. Gobbetti, J. A. I. Guitián, F. Marton, J. Nilsson, and G. Pintore, "An interactive 3d medical visualization system based on a light field display," *The Visual Computer*, vol. 25, 2009, to appear.