

Advances in massive model visualization in the CYBERSAR project

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

¹ CRS4, Pula, Italy, vic@crs4.it

Abstract—We provide a survey of the major results obtained within the CYBERSAR project in the area of massive data visualization. Despite the impressive improvements in graphics and computational hardware performance, interactive visualization of massive models still remains a challenging problem. To address this problem, we developed methods that exploit the programmability of latest generation graphics hardware, and combine coarse-grained multiresolution models, chunk-based data management with compression, incremental view-dependent level-of-detail selection, and visibility culling. The models that can be interactively rendered with our methods range from multi-gigabyte-sized datasets for general 3D meshes or scalar volumes, to terabyte-sized datasets in the restricted 2.5D case of digital terrain models. Such a performance enables novel ways of exploring massive datasets. In particular, we have demonstrated the capability of driving innovative light field displays able of giving multiple freely moving naked-eye viewers the illusion of seeing and manipulating massive 3D objects with continuous viewer-independent parallax.

I. INTRODUCTION

OVER the last several decades, there have been significant advances in model acquisition, computer-aided design (CAD), and simulation technologies. These technologies have resulted in massive data sets of complex geometric models. These data include complex CAD environments, natural landscape models, scanned urban data, and various scientific simulation data. These massive data typically require giga byte size and even tera byte size. Handling such massive models presents important challenges to software and system developers. This is particularly true for highly interactive 3D applications, such as visual simulations and

virtual environments, with their inherent focus on interactive, low-latency, and real-time processing. In the last decade, the graphics community has witnessed tremendous improvements in the performance and capabilities of computing and graphics hardware. Despite the continued increase in computing and graphics processing power, it is clear to the graphics community that one cannot just rely on hardware improvement to cope with arbitrarily large data sizes for the foreseeable future. This is not only because the increased computing power also allows users to produce more and more complex datasets, but also because memory bandwidth and data access speed grow at significantly slower rates than processing power and become the major bottlenecks when dealing with massive datasets. As a result, massive datasets cannot be interactively rendered by brute force methods. The real challenge is in designing rendering systems that capture as much of this performance growth as possible. When dealing with massive models, achieving high performance requires methods for carefully managing bandwidth requirements, controlling working set size, and ensuring coherent access patterns. Furthermore, given current multi-core CPUs and GPUs, solutions which can be formulated in a (data) parallel fashion will be able to continually take advantage of improved hardware performance. In this contribution, we provide a survey of the major results obtained within the project CYBERSAR in this field of research. Specifically, we will focus on the methods we developed to deal with huge models. The techniques exploit the programmability of latest generation graphics hardware, and combine coarse-grained multires-

olution models, chunk-based data management with compression, incremental view-dependent level-of-detail selection, and visibility culling. The models that can be interactively rendered with our methods range from multi-gigabyte-sized datasets for general 3D meshes or scalar volumes, to terabyte-sized datasets in the restricted 2.5D case of digital terrain models. Such a performance enables novel ways of exploring massive datasets. The rest of the paper is organized as follows: section II describes our multi-resolution out-of-core techniques for rapid high-quality visualization of textured digital terrain models and high resolution urban environments, section III deals with processing, distribution, and rendering of massive dense 3D meshes and huge complex 3D models, and section IV describes our method for rendering massive scalar volumes employing single pass GPU raycasting. Finally, section V describes our techniques for harnessing the power of novel 3D displays able of giving multiple freely moving naked-eye viewers the illusion of seeing and manipulating massive 3D objects with continuous viewer-independent parallax.

II. TERRAINS AND URBAN ENVIRONMENTS



Fig. 1. **Terrain rendering:** Interactive inspection of complex terrain datasets.

Real-time 3D exploration of remote digital elevation models built from high resolution imagery and elevation data has long been one of the most important components in a number of practical applications, and extensive research has been carried out in terms of methods and techniques for processing, distributing, and rendering very large datasets. The increased availability of broadband networks and high performance graphics PCs has made this technology, once limited to professional applications,

increasingly popular, as testified by the success of Internet geo-viewing tools like Google Earth, NASA WorldWind, and Microsoft Virtual Earth. The efficient implementation of such tools requires a combination of technologies for adaptively rendering high quality terrain views at high frame rates and techniques for efficiently streaming data from the systems serving the terrain database to a large number of rendering clients. In this field, we introduced the BDAM methods [1], multi-resolution out-of-core techniques for rapid high-quality visualization of textured digital terrain models and high resolution urban environments. The BDAM framework introduced one of the first methods for rapidly generating seamless variable resolution surfaces by assembling precomputed patches. In the context of the CYBERSAR project, we improved the BDAM framework by introducing a compressed multiresolution representation and a client-server architecture for supporting interactive high quality remote visualization of very large textured planar and spherical terrains[2]. The approach incrementally updates a chunked level-of-detail BDAM hierarchy by using precomputed wavelet coefficient matrices [3] decoded from a compressed bitstream originating from a thin server. The structure combines the aggressive compression rates of wavelet-based image representations with the ability to ensure overall geometric continuity for variable resolution views of planar and spherical terrains with no need for run-time stitching [4]. The efficiency of the approach was demonstrated on a large scale interactive remote visualization of global and local terrains on ADSL networks. A library implementing this approach [5] has been incorporated into widely distributed geo-viewing systems [6], [7] (see figure 1). A survey on semi-regular multiresolution models for interactive terrain rendering has been also presented [8], illustrating the principal techniques based on tiled blocks and nested regular grids, quadtrees and triangle bin-trees triangulations, as well as cluster based approaches.




Fig. 2. **Massive models:** Digital Michelangelo David 1mm Rendering.

III. MASSIVE MODELS

Many important applications domains, including 3D scanning, computer aided design, and numerical simulation, require the interactive inspection of huge geometric models. In this field, we considered methods for supporting inspection of surface models characterized by a high sample density, such as those generated by laser scanning. We also focused on methods able to support very large arbitrary surface models with high topological genus, highly variable depth complexity, fine geometric detail, "Bad" tessellations. Our main contributions are the introduction of Tetrapuzzles [9], a coarse grained multiresolution model based on hierarchical volumetric decomposition, that lead to the first GPU bound high quality technique for large scale meshes, and Far Voxels [10], an efficient approach for end-to-end out-of-core construction and interactive inspection that method tightly integrates visibility culling and out-of-core data management with a level-of-detail framework. The efficiency and generality of the approaches have been extensively demonstrated with the interactive rendering of extremely complex heterogeneous surface models on current commodity graphics platforms. The methods have been employed on a number of test cases, ranging from laser scans, including all Digital Michelangelo models, to isosurfaces, to extremely large CAD models, including the full Boeing 777 model (see figures 2 and 3).

Fully interactive performance is obtained even for large windows on current single processor PCs.

In the context of the CYBERSAR project, we introduced a cluster parallel multiresolution renderer able to drive a lightfield display [11], [12]. The method is a parallel spatial 3D display-aware version of the Adaptive Tetrapuzzles technique. It dynamically adapts model resolution by taking into account the particular spatial accuracy characteristics of the display.




Fig. 3. **Massive models:** Interactive inspection of complex datasets.

Between the most relevant results, surveys and contributions to specialized courses have been presented in the most important computer graphics conferences [13], [14], [15], [16], illustrating the massive model rendering methods to the graphics community.

IV. MASSIVE SCALAR VOLUMES




Fig. 4. **Massive scalar volumes rendering.** Interactive volume exploration of multi-giga voxels datasets on a desktop PC.

The ability to interactively render rectilinear scalar volumes containing billions of samples on desktop PCs is of primary importance for a number of applications, which include medical visualization and numerical simulation results analysis. In this field, in the context of the CYBERSAR project, we introduced MOVR [17], an adaptive out-of-core technique for rendering massive scalar datasets within a single-pass GPU raycasting framework. The method exploits an adaptive loader executing on the CPU for updating a working set of bricks maintained on GPU memory by asynchronously fetching data from an out-of-core volume octree representation. At each frame, a compact indexing structure, which spatially organizes the current working set into an octree hierarchy, is encoded in a small texture. This data structure is then exploited by an efficient raycasting algorithm, which computes the volume rendering integral by enumerating non-empty bricks in front-to-back order and adapting sampling density to brick resolution. In order to further optimize memory and bandwidth efficiency, the method also exploits feedback from the renderer to avoid refinement and data loading of occluded zones. The resulting method is extensible, fully adaptive, and able to interactively explore multi-giga-voxel datasets on a desktop PC (see figure 4).

V. INTERACTIVE VISUALIZATION ON NOVEL 3D DISPLAYS



Fig. 5. **Real-time interaction with the light field display.** Real-time inspection of a large volumetric model containing a CT dataset of a biological specimen. Pictures are taken with a hand-held camera at different viewing angles, in order to highlight the horizontal parallax of the light field display.

Multiscopic visualization is an emerging display technology that aims to reproduce three dimensional scenes by generating observer independent light fields. In the CYBERSAR project,



Fig. 6. **Massive model visualization on a 50Mpixel spatial 3D display.** The display has a screen dimension of 1.6x0.9 meters, and provides continuous horizontal parallax. Interactive rates in the manipulation of massive geometric models are guaranteed by a cluster-parallel spatial 3D display aware coarse grained multiresolution technique.

we contributed to this field by developing methods for driving displays of this type [11], [18]. Specifically, we built a parallel multiresolution rendering system [11] driving a spatial 3D display able to give multiple viewers the illusion of seeing virtual models floating at fixed physical locations situated in a human-scale working volume (see figure 6). Each viewer sees the scene from their point of view and enjoys full, continuous, horizontal parallax without specialized viewing devices. The efficiency of the approach is demonstrated by an application supporting interactive manipulation of massive colored highly tessellated models on a large (1.6x0.9 meters) 50Mpixel display that allows for a room-size working space. We also developed a prototype medical data visualization system exploiting a small scale light field display and custom direct volume rendering techniques to enhance understanding of massive volumetric data, such as CT, MRI, and PET scans [18], [19], [20]. The system can be integrated with standard medical image archives and extends the capabilities of current radiology workstations. The system allows multiple untracked naked-eye users in a sufficiently large interaction area to coherently perceive rendered volumes as real objects, with stereo and motion parallax cues (see figure 5). In this way, an effective collaborative analysis of volumetric data can be achieved. Evaluation tests demonstrated the usefulness of the generated

depth cues and the improved performance in understanding complex spatial structures with respect to standard techniques.

REFERENCES

- [1] P. Cignoni, F. Ganovelli, E. Gobbetti, F. Marton, F. Ponchio, and R. Scopigno, "BDAM: Batched dynamic adaptive meshes for high performance terrain visualization," *Computer Graphics Forum*, vol. 22, no. 3, pp. 505–514, Sept. 2003.
- [2] F. Bettio, E. Gobbetti, F. Marton, and G. Pintore, "High-quality networked terrain rendering from compressed bitstreams," in *Proc. ACM Web3D International Symposium*, ACM Press. New York, NY, USA, April 2007, pp. 37–44.
- [3] E. Gobbetti, F. Marton, P. Cignoni, M. D. Benedetto, and F. Ganovelli, "C-BDAM – compressed batched dynamic adaptive meshes for terrain rendering," *Computer Graphics Forum*, vol. 25, no. 3, September 2006, proc. Eurographics 2006.
- [4] P. Cignoni, F. Ganovelli, E. Gobbetti, F. Marton, F. Ponchio, and R. Scopigno, "Planet-sized batched dynamic adaptive meshes (P-BDAM)," in *IEEE Visualization*, 2003, pp. 147–154.
- [5] "Ratman," Url: ratman.sourceforge.net, 2009.
- [6] "Sardegna3d," Url: www.sardegna3d.it, 2007.
- [7] "Italia3d," Url: www.italia3d.it, 2008.
- [8] R. Pajarola and E. Gobbetti, "Survey on semi-regular multiresolution models for interactive terrain rendering," *The Visual Computer*, vol. 23, no. 8, pp. 583–605, 2007.
- [9] P. Cignoni, F. Ganovelli, E. Gobbetti, F. Marton, F. Ponchio, 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.
- [10] E. Gobbetti and F. Marton, "Far Voxels – a multiresolution framework for interactive rendering of huge complex 3d models on commodity graphics platforms," *ACM Transactions on Graphics*, vol. 24, no. 3, pp. 878–885, August 2005, proc. SIGGRAPH 2005.
- [11] F. Bettio, E. Gobbetti, F. Marton, and G. Pintore, "Scalable rendering of massive triangle meshes on light field displays," *Computers & Graphics*, vol. 32, no. 1, pp. 55–64, February 2008.
- [12] —, "Multiresolution visualization of massive models on a large spatial 3d display," in *Proc. Eurographics Symposium on Parallel Graphics and Visualization*, Eurographics Association. Aire-la-Ville, Switzerland, May 2007.
- [13] A. Dietrich, E. Gobbetti, and S.-E. Yoon, "Massive-model rendering techniques: A tutorial," *IEEE Computer Graphics and Applications*, vol. 27, no. 6, pp. 20–34, nov/dec 2007.
- [14] E. Gobbetti, D. Kasik, and S. eui Yoon, "Technical strategies for massive model visualization," in *Proc. ACM Solid and Physical Modeling Symposium*, ACM Press. New York, NY, USA, 2008, pp. 405–415. [Online]. Available: <http://www.crs4.it/vic/cgi-bin/bib-page.cgi?id=Gobbetti:2008:TSM>
- [15] S. eui Yoon, A. Dietrich, E. Gobbetti, D. Manocha, F. Marton, M. Agus, R. Pajarola, and P. Slusallek, "Interactive massive model rendering," SIGGRAPH Asia 2008 Course Notes, SIGGRAPH Asia 2008, Singapore, December 2008.
- [16] D. Kasik, A. Dietrich, E. Gobbetti, D. Manocha, F. Marton, P. Slusallek, A. Stephens, and S. eui Yoon, "Massive model visualization techniques," SIGGRAPH 2008 Course Notes, SIGGRAPH 2008, San Diego, August 2008.
- [17] 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.
- [18] 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.
- [19] 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.
- [20] 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.