Interactive Visualization at UFRGS:
Ongoing research at the Computer Graphics and Interaction Group

João Comba, Anderson Maciel, Carla Freitas,
Luciana Nedel, Manuel M. Oliveira
Instituto de Informática (INF) – Universidade Federal do Rio Grande do Sul (UFRGS)
Caixa Postal 15064, CEP 91501-970 Porto Alegre, Brazil – http://www.inf.ufrgs.br/cg

***Abstract*—The Computer Graphics and Interaction group at UFRGS has a well-established tradition of working in the subject of interactive visualization. With publications in the main venues of the field, such as the IEEE Visualization and Eurovis conferences, or journals such as IEEE TVCG (Transactions on Visualization and Computer Graphics) and Computer Graphics Forum (CGF), the group has consistently produced research in the area of interactive visualization throughout the years. Along the visualization group from USP-São Carlos, it organized the two previous instances of the workshop on interactive visualization, held at Sibgrapi 2007 and 2010. In this paper we summarize the recently published research in the period of 2010 to today. We also describe ongoing projects, that we hope can promote discussions and foster collaborations with other research groups.**

Keywords: Visualization; Interaction; Graphics; Information Visualization; Volume Visualization; High-Performance Computing

I. Published Research 2010-2012

In this section we review the work our lab published in the last two years in the area of interactive visualization.

A. Isocontouring and Feature Extraction of Higher-Order Data

Scientists and engineers are making increasingly use of hp-adaptive discretization methods to compute simulations. While techniques for isocontouring the high-order data generated by these methods have started to appear, they often do not allow interactive data exploration. In [1] we presented a novel interactive approach for approximating isocontouring of high-order data. Another way to evaluate higher-order surfaces is by extracting line-type features, such as creases (ridges and valleys). The parallel vectors (PV) operator is a feature extraction approach for defining such features in scalar fields, as well as separation, attachment, and vortex core lines in vector fields. In this work [2], [3] we extend PV feature extraction to higher-order data represented by piecewise analytical functions defined over grid cells. The extraction uses PV in two distinct stages. First, seed points on the feature lines are placed by evaluating the inclusion form of the PV criterion with reduced affine arithmetic. Second, a feature flow field is derived from the higher-order PV expression where the features can be extracted as streamlines starting at the seeds. Our approach allows for bounds regarding accuracy with respect to existence, position, and topology of the features obtained (Figure 1).

B. Visualization of the Human TMJ Behavior

The temporomandibular joint (TMJ) is one of the most important and complex joints of the body and its pathologies affect a

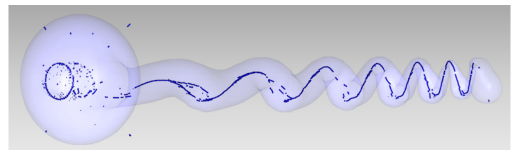

Fig. 1. ***Efficient Parallel Vectors Feature Extraction from Higher-Order Data***: Valley lines extracted from a Discontinuous Galerkin simulation.

great percentage of the human population. The simulation and real-time visualization of the TMJ behavior during opening, closing and chewing movements can be very useful to the understanding of this articulation by physicians, helping them to prevent or fix problems due to accidents or diseases. In this work we proposed a model to simulate the human TMJ behavior based on the concept of two interdependent joints [4]. The model was conceived using multimodal information acquired from CT and MRI images of a live person, as well as motion data acquired from this same person with a magnetic motion capture device (Figure 2). Simulation of movement of other TMJs – based on different morphology of bones and teeth – were obtained by adapting the regular captured motion data through collision detection and treatment methods. The model was evaluated with image registration techniques by comparing simulated results with real, captured motion data. We also validated how the model can predict TMJ behavior in the presence of normal or abnormal bones and teeth morphologies.

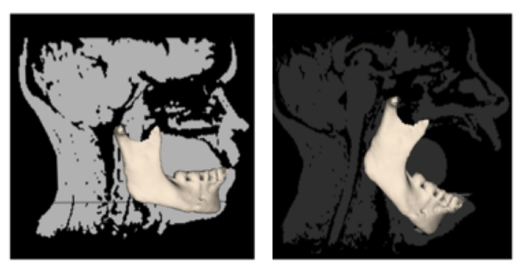

Fig. 2. Lateral views of the registration between the jaw region segmented on MRI images, and the corresponding mesh extracted from the TMJ model.

C. Parallel Dataflow Frameworks for Interactive Visualization

Dataflow systems play a central role in the description of computational pipelines in visualization systems. Since pipelines can be complex and composed of several branches, the efficient computation leveraging parallel computation is an active area of research. In [5], [6] we propose a new framework design for exploiting multi-core architectures in the context of visualization dataflow systems. In [7] we proposed a new dataflow architecture, called *HyperFlow*, that offers a supporting infrastructure that creates an abstraction layer over computation resources and naturally exposes heterogeneous computation (e.g. CPUs and GPUs) to dataflow processing. We have included a set of synthetic and real-case applications to show the efficiency of our system. We designed a general suite of micro-benchmarks that captures main parallel pipeline structures and allows evaluation of HyperFlow under different stress conditions. We demonstrated the potential of our system with relevant applications in visualization. Implementations in HyperFlow have greater performance than actual hand-tuning codes, yet still providing high scalability on different platforms.

D. Parallel visualization on large clusters using MapReduce

Exploratory visualization systems are increasingly expected to support scalable data manipulation, restructuring, and querying capabilities in addition to core visualization algorithms. We posit that new emerging abstractions for parallel data processing, in particular computing clouds, can be leveraged to support large-scale data exploration through visualization. In [8] we take a first step into evaluating the suitability of the MapReduce framework to implement large-scale visualization techniques. MapReduce is a lightweight, scalable, general-purpose parallel data processing framework increasingly popular in the context of cloud computing. Specifically, we implement and evaluate a representative suite of visualization tasks (mesh rendering, isosurface extraction, and mesh simplification) as MapReduce programs, and report quantitative performance results applying these algorithms to realistic datasets (Figure 3).

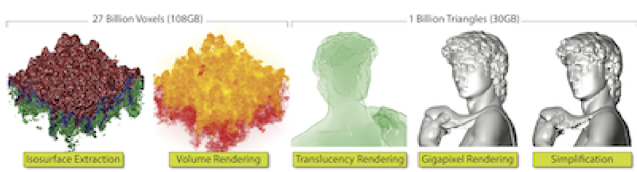

Fig. 3. ***Parallel Visualization on Large Clusters using MapReduce***: Set of visualization tasks implemented using *Map/Reduce*: isosurface extraction, volume rendering, and multi-resolution mesh rendering.

E. Parallel Large Data Visualization with Display Walls

In recent years, scientific data has become large and complex enough to require specifically designed software tools for visualization. Moreover, the amount and complexity of different visualization techniques available introduce a bottleneck on the effectiveness of the data exploration process. As datasets increase in complexity, single-display systems become progressively less effective. Tiled displays have gained popularity as easy to build low-cost devices which can extend available screen space. They can be constructed in several ways, but we restricted ourselves to arrays of high-resolution LCD panels. In [9] we presented PVW (Parallel Visualization using display Walls), a framework that uses display walls for scientific visualization, requiring minimum labor in setup, programming and configuration. PVW works as a plug-in to pipeline-based visualization software, and allows users to migrate existing visualizations designed for a single workstation, single-display setup to a tiled display on a distributed machine (Figure 4).

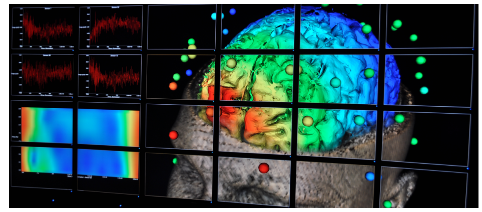

Fig. 4. ***Parallel Large Data Visualization with Display Walls***: Display walls can be used to inspect complex data in high resolution.

F. Interaction with display walls using levels of precision

The use of large displays are giving rise to new working scenarios in which users are not sitting in front of the screen nor they have a table upon which to pose their mice. A widespread example of this situation are the multi-display screens in operation centers, public spaces and scientific facilities. While they offer great visualization possibilities, efficient interaction with such displays is still a major challenge.

In order to fullfil this gap, we propose levels of precision (LOP) cursor [10], a metaphor for high precision pointing and simultaneous cursor controlling using commodity mobile devices. The LOP-cursor uses a two levels of precision representation that can be combined to access low and high resolution of input. It provides a constrained area of high resolution input and a broader area of lower input resolution, offering the possibility of working with a two legs cursor using only one hand. LOP-cursor is designed for interaction with large high resolution displays, e.g. display walls, and distributed screens/computers scenarios. Targets smaller than 0.3 cm can be selected by users at distances over 1.5 m from the screen with minimum effort. We are now integrating the LOP-cursor with a social network visualization tool and will evaluate its behavior with dense graphs.

G. Image-Set Processing Streaming Architecture

Atlas construction is an important technique in medical image analysis that plays a central role in understanding the variability of brain anatomy. The construction often requires applying image processing operations to multiple images (often hundreds of volumetric datasets), which is challenging in computational power as well as memory requirements. In [11], [12] we introduced MIP, a Multi-Image Processing streaming framework to harness the processing power of heterogeneous CPU/GPU systems. MIP is composed of specially designed streaming algorithms and data structures that provides an optimal solution for out-of-core multi-image processing problems both in terms of memory usage and computational efficiency. MIP makes use of the asynchronous execution mechanism supported by parallel heterogeneous systems to efficiently hide the inherent latency of the processing pipeline of out-of-core approaches. We demonstrated the efficiency of the MIP framework on synthetic and real datasets.

H. Color-Encoding Patterns

Color is one of the most efficient ways of conveying information in visualization applications. Color vision deficiency (CVD) affects approximately 200 million individuals around the world, considerably affecting their ability to perform color-related tasks. We have proposed a technique for using patterns to encode color information for individuals with CVD, in particular for dichromats [13]. This approach is the first content-independent method to overlay patterns on colored visualization contents that not only minimizes ambiguities but also allows color identification. Moreover, since overlaying patterns does not compromise the underlying original colors, it does not impair the perception of normal trichromats. We have validated this approach with two user studies: the first one included 11 subjects with CVD and 19 normal trichromats, and focused on images that use colors to represent multiple categories; the second, included 16 subjects with CVD and 22 normal trichromats, and considered a broader set of images. The results of these studies show that overlaying patterns significantly improves the performance of dichromats in several color-based visualization tasks, making their scores similar to normal trichromats?. More interestingly, the use of overlaid patterns augments color information in a positive manner, providing normal trichromats with a way of achieving finer color discrimination (Figure 5).

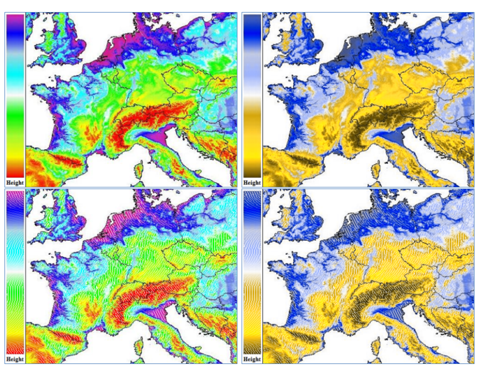

Fig. 5. Using pattern to encode color information. Original image (top left), and how it is perceived by an individual with deuteranopia (CVD) (top right). Original image with superimposed patterns (bottom left), and how it perceived by an individual with deuteranopia (bottom right).

II. Ongoing Projects

In this section we discuss ongoing work being developed in the lab, to promote discussion and foster collaborations.

A. Heart Rate Visualization of Multiple Running Activities

Heart rate (HR) monitors are very convenient devices for measuring body statistics during physical activities, such as heartbeat frequency, speed, cadence, geophysical data, etc. Such data can be composed of simple statistics about the exercise (e.g. mininum, average and maximum values of a given quantity) to the full time-series data collected at every second of the exercise. The recent widespread availability of heart rate monitors is allowing people with different physical fitness levels to measure activity during and after exercise. This can be extremely important for people with low fitness levels, since they are susceptible to cardiovascular diseases or other physical injuries when exercising at high heartbeat frequencies. In particular, the post-exercise visualization of the time-series data of heart rate measurements allows a more sophisticated analysis to be performed, and most monitors provide tools to export and display this information for a given exercise. However, the ability to display multiple activities is limited or non-existent for many of the existing software. In this research we are studying ways to perform a visual analysis of HR time-series data from multiple activities.

B. Visualizing Parallel Dataflow Execution Traces

Most performance analysis tools focus on presenting an overload of details, with little application-dependent structure, and predefined statistical summaries. This makes the complex relations present in a parallel program not directly recognizable to the user, making the task of identifying and dealing with performance issues more costly in both time and effort. In this line of research we are currently investigating the requirements to create visualizations of execution traces of parallel programs modeled as dataflows. We propose the Smart Trace concept, to encode the structure present in the data, and guide the construction of specialized visualizations. A visualization tool can then leverage the relationships in the data to automate a series of visualization tasks. We show with examples the power and flexibility of visualizations we can create to address specific questions formulated about the analysis of the data, with emphasis in parallel dataflow traces.

C. VR- and AR-based tools for planning medical interventions

We propose new interactive visualization tools for analyzing 3D medical data and to support diagnosis and treatment planning. Our first case studies are on anatomic hepatectomies (Fig. 6). Important decision making information, as organ and tumor volumes, or estimation of remaining functional tissue after surgery, are obtained from direct interaction on patient specific 3D data (CT or MRI). Part of the visualization is done using mobile display devices, allowing augmented reality views directly on the patient's body [14]. Figure 7 shows a tablet-based visualization for CT. Our current research focuses on extending the approach to multiple image modalities. In such a way, 2D real time imaging could be combined with detailed previously acquired 3D data. This also includes non-rigid registration of such data through physics-based tissue deformation.

D. Interactive visualization of social networks

Social networks have been studied along the years but only recently they have become used as a mean of communication.

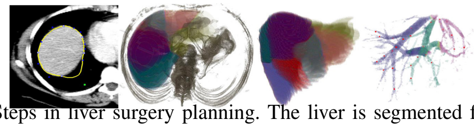

Fig. 6. Steps in liver surgery planning. The liver is segmented from CT and functional segments are classified from vessel branching directly on the 3D volume.

Fig. 7. An augmented reality tool for interactive visualization of the volume generated from CT. It allows the analysis of inner body structures by pointing the mobile display directly upon the body.

The group has started to work with social networks as an experimental area for graph visualization techniques. A technique (and tool) named MagnetViz was designed for the interactive manipulation of force-directed graph layouts, allowing users to obtain visualizations based on the graph topology and/or the attributes of its nodes and edges [15]. Users can introduce virtual magnets anywhere in the layout and these can be set to attract nodes and edges that fulfill user-defined criteria. Scientific collaboration (co-authorship) networks are being used as main datasets for experiments.

E. Interactive volume visualization and analysis

Many applications deal with volumetric data, which can be visualized either directly, using transfer functions that map data values to color and opacity values, or as surfaces that are extracted from the volume by some thresholding process. Both categories of visualization algorithms have been investigated in successive projects within our group as well as ways of interacting with volumes either for measuring inner structures or carving or adding material as a sculptor. Recent results allow interactive visualization of the volume interior displayed as illustrations to enhance perception of features [16]. An ongoing project aims at analyzing structures in images of liver samples taken from a confocal microscope.

III. Final Remarks

As seen in the previous sections, projects in interactive visualization range from scientific and information visualization applications, with a great interest on performance enhancement, with focus on high performance computing and parallel architectures. This diversity provides an interesting work environment, fostering collaboration between students. Due to these characteristics, the group welcomes MSc and PhD students with varied background but with strong evidence of a research-oriented profile, and also seeks collaboration with other groups to broaden research opportunities.

The full list of group's publications can be found at www.inf.ufrgs.br/cg.

References

[1] C. Pagot, J. Vollrath, F. Sadlo, D. Weiskopf, T. Ertl, and J. ao Luiz Dihl Comba, "Interactive Isocontouring of High-Order Surfaces," in *Scientific Visualization: Interactions, Features, Metaphors*, ser. Dagstuhl Follow-Ups, H. Hagen, Ed. Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2011, vol. 2, pp. 276–291. [Online]. Available: http://drops.dagstuhl.de/opus/volltexte/2011/3305

[2] C. Pagot, D. Osmari, F. Sadlo, D. Weiskopf, T. Ertl, and J. Comba, "Efficient parallel vectors feature extraction from higher-order data," *Computer Graphics Forum*, vol. 30, no. 3, pp. 751–760, 2011. [Online]. Available: http://dx.doi.org/10.1111/j.1467-8659.2011.01924.x

[3] F. Sadlo, M. andffinger, C. Pagot, D. Osmari, J. Comba, T. Ertl, C.-D. Munz, and D. Weiskopf, "Visualization of cell-based higher-order fields," *Computing in Science Engineering*, vol. 13, no. 3, pp. 84 –91, may-june 2011.

[4] M. B. Villamil, L. P. Nedel, C. M. D. S. Freitas, and B. Macq, "Simulation of the human tmj behavior based on interdependent joints topology," *Comput. Methods Prog. Biomed.*, vol. 105, no. 3, pp. 217–232, Mar. 2012. [Online]. Available: http://dx.doi.org/10.1016/j.cmpb.2011.09.010

[5] H. T. Vo, D. K. Osmari, B. Summa, J. L. D. Comba, V. Pascucci, and C. T. Silva, "Streaming-enabled parallel dataflow architecture for multicore systems," *Computer Graphics Forum*, vol. 29, no. 3, pp. 1073–1082, 2010. [Online]. Available: http://dx.doi.org/10.1111/j.1467-8659.2009.01704.x

[6] H. Vo, J. Comba, B. Geveci, and C. Silva, "Streaming-enabled parallel data flow framework in the visualization toolkit," *Computing in Science Engineering*, vol. 13, no. 5, pp. 72 –83, sept.-oct. 2011.

[7] H. T. Vo, D. K. Osmari, J. Comba, P. Lindstrom, and C. T. Silva, "Hyperflow: A heterogeneous dataflow architecture," in *EGPGV*, 2012, pp. 1–10.

[8] H. Vo, J. Bronson, B. Summa, J. Comba, J. Freire, B. Howe, V. Pascucci, and C. Silva, "Parallel visualization on large clusters using mapreduce," in *Large Data Analysis and Visualization (LDAV), 2011 IEEE Symposium on*, oct. 2011, pp. 81 –88.

[9] L. Scheidegger, H. T. Vo, J. Kruger, C. T. Silva, and J. L. D. Comba, "Parallel large data visualization with display walls," in *Visualization and Data Analysis 2012*, P. C. Wong, D. L. Kao, M. C. Hao, C. Chen, R. Kosara, M. A. Livingston, J. Park, and I. Roberts, Eds., vol. 8294, no. 1. SPIE, 2012, p. 82940C. [Online]. Available: http://link.aip.org/link/?PSI/8294/82940C/1

[10] H. Debarba, L. P. Nedel, and A. Maciel, "Lop-cursor: Fast and precise interaction with tiled displays using one hand and levels of precision." in *3DUI*, M. Billinghurst, J. J. L. Jr., and A. Lcuyer, Eds. IEEE, 2012, pp. 125–132.

[11] L. K. Ha, J. Krüger, J. Comba, S. C. Joshi, and C. T. Silva, "Optimal multi-image processing streaming framework on parallel heterogeneous systems," in *EGPGV*, 2011, pp. 1–10.

[12] L. K. Ha, J. Kruger, J. L. D. Comba, C. T. Silva, and S. Joshi, "Isp: An optimal out-of-core image-set processing streaming architecture for parallel heterogeneous systems," *Visualization and Computer Graphics, IEEE Transactions on*, vol. 18, no. 6, pp. 838 –851, june 2012.

[13] B. Sajadi, A. Majumder, M. M. Oliveira, R. G. Schneider, and R. Raskar, "Using patterns to encode color information for dichromats," *IEEE Transactions on Visualization and Computer Graphics*, vol. 99, no. PrePrints, 2012.

[14] H. Debarba, J. Grandi, D. Zanchet, and A. Maciel, "Anatomic hepatectomy planning throug mobile display visualization and interaction," in *MMVR19: Medicine Meets Virtual Reality 19*, ser. Studies in Health Technology and Informatics, J. D. Westwood, R. S. Haluck, H. M. Hoffman, G. T. Mogel, R. Phillips, R. A. Robb, and K. G. Vosburgh, Eds., vol. 173. IOS Press, 2012, pp. 111–115.

[15] A. S. Spritzer and C. M. D. S. Freitas, "Design and evaluation of magnetviz: A graph visualization tool," *IEEE Transactions on Visualization and Computer Graphics*, vol. 18, pp. 822–835, 2012.

[16] F. de Moura Pinto and C. M. D. S. Freitas, "Illustrating volume data sets and layered models with importance-aware composition," *The Visual Computer*, vol. 27, no. 10, pp. 875–886, 2011.