

SAGE: the Scalable Adaptive Graphics Environment

RENAMBOT, Luc (luc@evl.uic.edu), RAO, Arun (arao@evl.uic.edu), SINGH, Rajvikram (rsingh@evl.uic.edu), BYUNGIL, Jeong (bijeong@evl.uic.edu), KRISHNAPRASAD, Naveen (naveen@evl.uic.edu), VISHWANATH, Venkatram (venkat@evl.uic.edu), CHANDRASEKHAR Vaidya (vchand2@evl.uic.edu), SCHWARZ Nicholas (schwarz@evl.uic.edu), SPALE, Allan (aspale@evl.uic.edu), ZHANG Charles (charles@evl.uic.edu), GOLDMAN Gideon (ggoldm2@uic.edu), LEIGH, Jason (spiff@evl.uic.edu), JOHNSON, Andrew (aej@evl.uic.edu)

Electronic Visualization Laboratory, Dept. of Computer Science,
University of Illinois at Chicago

1. Introduction

We envision situation-rooms and research laboratories in which all the walls are made from seamless ultra-high-resolution displays fed by data streamed over ultra-high-speed networks, from distantly located visualization and storage servers [1]. It will allow local and distributed groups of researchers to work together on large amounts of distributed heterogeneous datasets. From our prior work on the Continuum [2], we have learned that it is crucial for collaborators to have both local control (e.g. on a tablet or laptop) and the casual ability to share their work and see what others are working on (e.g. on a large tile display). We are taking the next steps toward this vision by building *LambdaVision* - an 11x5 tiled display with a total resolution of 100 megapixels and developing SAGE, the Scalable Adaptive Graphics Environment (see Figure 1). SAGE will allow the seamless display of various networked applications over the whole display. Each visualization application (such as 3D rendering, remote desktop, video streams, 2D maps) streams its rendered pixels (or primitives) to a virtual high-resolution frame buffer, allowing for any given layout onto the display.




Figure 1 - Example of a collaborative SAGE session with multiple applications running simultaneously

Computing platforms typically provide some notion of remote desktop access. Tools such as VNC or Microsoft's Remote Desktop were designed to transmit screens of single desktops to

remote computers over slow networks. They are designed to operate on event triggered streaming mechanisms that are not suitable for real-time streaming of scientific visualization or collaborative applications. Flexible scalable graphics systems such as Chromium [7] and Aura [11] are designed for distributing visualization to and from cluster driven tiled-displays. Applications have to be specifically written with the Aura API to enable them to transmit their graphics and Chromium was not designed to stream graphics over wide-area networks. IBM's Scalable Graphics Engine [9] is also a hardware-based approach that allows for reception of pixels streamed over networks and can drive an array of displays synchronously. Currently, it is limited to a 16-port output and does not support protocols for streaming over wide-area networks. In all these approaches, it is difficult for users to add their own video sources, compression modules or network protocols to the systems, which are needed for high-bandwidth wide-area networks.

2. Design

The impetus for the work on SAGE comes from the two main areas. The first is that applications written for one graphics environment have to be re-designed before they can be run under other environments. For example, visualization tools that are developed for desktop computers are rarely able to take advantage of the processing power of a cluster of graphics computers; conversely visualization tools developed for clusters rarely function on desktop computers. Secondly, the ability of visualization software and systems to scale in terms of the amount of data they can visualize, and the resolution of the desired visualization, is still an area of intensive computer graphics research [5]. SAGE addresses the need to support heterogeneity and scalability by decoupling graphics rendering from graphics display and taking advantage of affordable ultra-high-bandwidth networks to bridge them.



GeoWall-2 demonstrated at Supercomputing 2003, showing the ocean floor bathymetry of the entire earth, data courtesy of USGS.



55-Tile LambdaVision display

Figure 2 – GeoWall-2 and LambdaVision displays

The SAGE architecture consists of a number of rendering resources (from single desktop computers to clusters of local or distributed PCs capable of rendering graphics either with dedicated graphics hardware or software), connected over the network to a scalable frame buffer, in our case the *LambdaVision* display (see Figure 2). *LambdaVision* pushes the concept of tiled-display technology [4,5,8] and anticipates the geosciences community's needs by coupling massive data storage (50TB of local storage) to a wall of 55 screens driven by 64-bit

computers attached to each other by tens-of-gigabits of networking, for a total over 100Mpixels. The SAGE environment with visualization software, network controls, data management, and human-computer interfaces will allow collaborative work for scientists. LambdaVision is developed after the GeoWall-2, which is a tiled display built from an array of 5x3 LCD tiles and driven by a cluster of nine, 32-bit PCs (each PC driving two displays, and one PC as the master controller). GeoWall-2 was demonstrated at both the Supercomputing (SC2003) and AGU conferences where various institutions displayed their datasets using this system.

Each application running in SAGE will stream imagery using the SAGE Application Interface Library (SAIL) - a highly optimized software library for lossless streaming of pixels over high-speed networks, based on TeraVision [3], our prior work in high performance graphics digitization and high-speed data streaming over gigabit networks. This decoupling of the rendering process from the display allows SAGE to scalably support distance collaboration by streaming the pixels to all the participating endpoints using either traditional router-based multicasting or photonic multicasting. Different endpoints in the collaboration may be seeing the same visualization on different devices with different display characteristics.




Figure 3 - SAGE Distributed Resource Allocation and Management

a. Visualization

The main components of the visualization system are designed to display high-resolution imagery, geo-referenced datasets (topography) and large 3D volumes. This work depends on some of our on-going research, the JuxtaView and Vol-a-Tile applications.

Visualization jobs invoked by a user are automatically dispatched to the optimal visualization resource currently available on the network. The chosen resource performs the visualization and streams the resulting imagery for display (see Figure 3). Note that, in some cases, it is more economical to render the visualization directly on the tiled display since the data is small enough to fit in the graphics cards' memory. In the first implementation, the only data streamed from the application are raw pixels, but we envision supporting "smart-pixels" or primitives (such as pixels plus depth information, layers of pixels, polygons in display lists...) to enable more efficient rendering.

JuxtaView [4] is a cluster-based application for viewing ultra-high-resolution images on scalable tiled displays. JuxtaView presents a new parallel computing and distributed memory approach for out-of-core montage visualization, using LambdaRAM, a software-based network-level cache system. JuxtaView enables a user to interactively roam through potentially terabytes of distributed, spatially referenced image data sets such as those generated from electron microscopes, satellites and aerial photography. The data is distributed using LambdaRAM, on the memory of all nodes of a PC cluster driving a tiled display wall or on a remote cluster. Aggressive pre-fetching schemes employed by LambdaRAM help to reduce latency involved in remote memory access. This approach outperforms by at least a factor of two a traditional memory-mapped file approach for out-of-core visualization.

Vol-a-Tile [14] is an interactive tool for exploring large volumetric data on scalable tiled displays. Hardware texture mapping and level-of-detail techniques provide interactivity. OptiStore is the data management server that provides the high performance I/O needed to stream data from storage to the nodes driving the tiled display over high-bandwidth photonic networks. It is designed to handle common data management operations, including loading data, maintaining meta-information about the data, such as dimension and scaling, and storing the results of sampled and cropped data. In addition, OptiStore provides data processing capabilities, including run time gradient/histogram generation, sampling, and cropping by utilizing functionality in VTK. Network connectivity is provided through the QUANTA network toolkit, which utilizes the aggressive network protocol, Reliable Blast UDP [10].

b. Networking

Such demanding visualization applications need large amount of bandwidth to access the datasets stored remotely. LambdaRAM [13] is the middleware designed based on our networking expertise designated to help OptiPuter [1] applications harness high bandwidth network connectivity and address latency of long haul networks. LambdaRAM achieves this by managing distributed memory pools for data storage and employing aggressive pre-fetching schemes to cache data locally for applications. LambdaRAM is based on the concept of Network Memory (NetRAM). The high bandwidth photonic networks interconnecting components on the OptiPuter makes NetRAM over wide areas practical. The concept behind NetRAM is to provide a massive pool of physical memory that is distributed over separate computers that are interconnected using high-speed photonic networks. In scenarios where datasets are too large to be stored locally, using NetRAM to move data between sites with high bandwidth connectivity becomes an attractive option. LambdaRAM, our NetRAM implementation, provides a read-only data cache, which fits most visualization and data mining applications where the original data is rarely modified. Application developers using LambdaRAM can specify the partitioning of data on distributed workstations to maximize locality of reference. LambdaRAM can overcome the latency involved in long distance photonic networks limited by the speed of light, through aggressive pre-fetching schemes. LambdaRAM is the network middleware that enables interactive visualization applications in the SAGE environment and on the LambdaVision display.

c. Collaboration

Collaborative work is enabled at two levels: either through the sharing of datasets using LambdaRAM or through the resulting pixels generated by the visualization application and distributed by the TeraVision package (in software or hardware mode).

TeraVision is a scalable platform-independent solution that is capable of transmitting multiple synchronized high-resolution video streams between single workstations and/or clusters without

requiring any modifications to be made to the source or destination machines. TeraVision can synchronize individual video streams to form a single larger stream; to scale and route streams generated by an array of MxN nodes to fit an XxY display; and exploits a variety of transport protocols such as reliable broadcast UDP and IP multicast. We conducted experiments over gigabit local-area networks and wide-area, streaming a 1600x1200 desktop at 10 frames per second or a 1024x768 laptop at more than 30 frames per second.

The SAGE architecture will allow multiple rendering nodes or clusters to access a virtual framebuffer across the network. The framework will intelligently partition the graphics pipeline to distribute the load. Factors such as computing and rendering capabilities of the participating machines will decide the load distribution. Thus, unlike TeraVision where only pixels are routed to the display machines, SAGE would also route geometry and custom graphics formats. The framework will also support the notion of multiple collaborators simultaneously accessing a display space through a shared “window” manager.

SAGE is designed to support data fusion for very large data sets. Collaborators at multiple remote sites with heterogeneous displays will be able to share and simultaneously visualize multiple datasets. For instance, the users could be simultaneously viewing high-resolution aerial or satellite imagery, as well as volumetric information on earthquakes and ground water. Audio and video conferencing between the groups will be handled over the AccessGrid, which streams are treated as another form of data that can be displayed on LambdaVision.

d. Interaction

The collaborators interact with SAGE’s FreeSpace Manager (see Figure 3), which provides an intuitive interface for moving visualizations around the tiled display. When a visualization window is moved from one portion of the screen to another, the FreeSpace Manager informs the remote rendering resource of the new destination for the streamed pixels, giving the user the illusion that they are working on one continuous computer screen, even though the systems performing the visualizations may be several thousand miles apart. The FreeSpace Manager is akin to a traditional desktop manager in a windowing system, except that it can scale from a single tabletPC screen to a desktop spanning 100 million pixels.

Users will control the SAGE applications from tablet PCs, laptops, or more exotic interfaces such as the *LambdaTable* – a table made of LCD screens with camera tracking. This decoupling of the user interface from the visualization allows users to easily move control of a given application between users at the same site or remote sites. Tablet PCs also give the users an easy ability to highlight and annotate information for the rest of the group. This can be especially important in areas such as annotating high-resolution core samples taken from several hundred feet below the Earth’s surface.

3. Implementation

The current implementation is in its early stage. We are building on the TeraVision experience, and have implemented an initial version of the FreeSpace Manager and the SAIL library for OpenGL applications. Figure 4 shows SAGE running on a 15-tile GeoWall-2 display: it depicts a sequential application (standard ‘atlantis’ GLUT demo) running on one node of the cluster, and the resulting pixels of the application being streaming to 12 nodes of the display. The OpenGL application requires minimal modifications: a few lines of code are added to link to the SAIL library. Note that the borders of the display are taken into account. The next steps are to support multiple applications, the ability to move and scale application windows across the display, and to enable more application through the SAIL library (VNC, AccessGrid videos, hardware capture with TeraVision). We are also actively developing the next version of high-bandwidth UDP-

based protocol, so called LambdaStream, which will offer high-bandwidth, low latency and low jitter over long fat networks. Such functionalities are required for networked visualization applications. The progress of the SAGE project can be monitored on the following web site: www.evl.uic.edu/cavern/sage




Figure 4 - Test application running in the SAGE environment

References

1. Larry L. Smarr, Andrew A. Chien, Tom DeFanti, Jason Leigh, Philip M. Papadopoulos "The OptIPuter" Communications of the ACM, Volume 46, Issue 11, November 2003, pp. 58-67.
2. Park, K., Renambot, L., Leigh, J. and Johnson, A., "The Impact of Display-rich Environments for Enhancing Task Parallelism and Group Awareness in Advanced Collaborative Environments", In Workshop on Advanced Collaborative Environments, June 22-24, 2003, Seattle, WA.
3. Rajvikram Singh, Byungil Jeong, Luc Renambot, Andrew Johnson and Jason Leigh "TeraVision: a Distributed, Scalable, High Resolution Graphics Streaming System", in the proceedings of IEEE Cluster 2004, San Diego, September 20-23, 2004.
4. Naveen K. Krishnaprasad, Venkatram Vishwanath, Shalini Venkataraman, Arun G. Rao, Luc Renambot, Jason Leigh, Andrew E. Johnson, and Brian Davis "JuxtaView - a Tool for Interactive Visualization of Large Imagery on Scalable Tiled Displays", in the proceedings of IEEE Cluster 2004, San Diego, September 20-23, 2004.
5. W. Blanke, C. Bajaj, D. Fussell, and X. Zhang, "The Metabuffer: a Scalable Multiresolution Multidisplay 3-D Graphics System using Commodity Rendering Engines." Tr2000-16, University of Texas at Austin, February 2000.
6. J. Leigh, L. Renambot, T.A. DeFanti, M.D. Brown, E. He, N.K. Krishnaprasad, J. Meerasa, A. Nayak, K. Park, R. Singh, S. Venkataraman, C. Zhang, D. Livingston, M. McLaughlin, "An Experimental OptIPuter Architecture for Data-Intensive Collaborative Visualization", 3rd Workshop on Advanced Collaborative Environments, Seattle, WA, June 2003
7. G. Humphreys, M. Houston, Y. Ng, R. Frank, S. Ahern, P. Kirchner, and J. T. Klosowski, "Chromium: A Stream-Processing Framework for Interactive Rendering on Clusters", Proceedings of SIGGRAPH 2002.
8. Y. Chen, H. Chen, D. W. Clark, Z. Liu, G. Wallace, K. Li, "Software environments for cluster-based display systems", First IEEE/ACM International Symposium on Cluster Computing and the Grid, May 2001.
9. Kenneth A. Perine, Donald R. Jones, William R. Wiley "Parallel Graphics and Interactivity with the Scaleable Graphics Engine" Proceedings of the 2001 ACM/IEEE conference on Supercomputing.
10. E. He, J. Leigh, O. Yu, T. A. DeFanti, "Reliable Blast UDP: Predictable High Performance Bulk Data Transfer", IEEE Cluster Computing 2002, Chicago, IL, Sept 2002.
11. D. Germans, H.J.W. Spoelder, L. Renambot, H. E. Bal, "VIRPI: A High-Level Toolkit for Interactive Scientific Visualization in Virtual Reality", Proc. Immersive Projection Technology/Eurographics Virtual Environments Workshop, Stuttgart, May 2001.
12. C. Zhang, J. Leigh, T. A. DeFanti, "TeraScope: Distributed Visual Data Mining of Terascale Data Sets over Photonic Networks", Journal of Future Generation Computer Systems (FGCS), Elsevier Science Press, Volume 19-6, August 2003
13. Eric He, Javid Alimohideen, Josh Eliason, Naveen Krishnaprasad, Jason Leigh, Oliver Yu, Thomas A. DeFanti, "Quanta: A Toolkit for High Performance Data Delivery over Photonic Networks," Journal of Future Generation Computer Systems (FGCS), Elsevier Science Press, Volume 19-6, August 2003.
14. Nicholas Schwarz, Shalini Venkataraman, Luc Renambot, Naveen Krishnaprasad, Venkatram Vishwanath, Jason Leigh, Andrew Johnson, Graham Kent, Atul Nayak, "Vol-a-Tile - a Tool for Interactive Exploration of Large Volumetric Data on Scalable Tiled Displays", IEEE Visualization 2004, Poster session, Austin, TX, October 2004.