Photogrammetric Week '07 Dieter Fritsch (Ed.) Wichmann Verlag, Heidelberg, 2007

Towards Gigapixel Display for Data Visualization

THOMAS ERTL, Stuttgart

EXTENDED ABSTRACT

The rapid development of digital camera technology has increased image resolution by an order of magnitude in the recent years with 20 megapixel SLR cameras now being widely available. Even gigapixel images are now made possible by special cameras or scanners or by stitching together hundreds of images taken with automatic mounts. However, display technology has not made similar progress. In television sets we are still waiting for the move from PAL (0.5 megapixels) to HDTV (2 megapixels) and computer monitors range from a typical 1 megapixels to 4 megapixels on the newer 30 inch wide screen LCDs. Projector technology is in the same range from 1.5 megapixels for the 1400 x 1050 customer beamers to the still unique 4K x 2K bulky LCOS cinema projectors. In order to display gigapixel imagery one can either pan and zoom such images on traditional displays or one can reverse the stitching idea and try to build tiled displays composed of several single display surfaces.

The first approach gained some attention at this year's ACM SIGGRAPH conference when researchers from the Universität Konstanz and Microsoft Research demonstrated how gigapixel images can be stitched together and how they can be viewed in a web browser plugin which loads image tiles from a pyramid representation, adapts the projection according to the field-of-view and compensates for high dynamic range differences by tone mapping. The second approach falls into two classes: tiled LCD walls and projector arrays. The largest display so far, the LambdaVision system at EVL, is a wall of 5 x 11 LCD panels with about 100 megapixels spread over a width of 5 meters. The main disadvantages of these LCD walls are the clearly visible borders (mullions) between tiles which give the impression of looking through a fence onto an image behind. Projector arrays have the potential of seamless tiling which, however, require sophisticated calibration of geometry and luminance. Another advantage of rear-projection walls is the potential for displaying stereo image pairs either passively by doubling the number of projectors or by active stereo with DLP projectors with a 120 Hertz refresh rate. The currently most advanced projection system is the LANL La Cueva Grande, a five-sided cave with 33 projectors and 43 megapixels in total.

Since current high-end graphics cards can drive displays up to a resolution of at most 2560 x 1600 pixels on 2 DVI outputs, each display larger than that needs to be connected to multiple graphics cards. A single modern PC system can hold at most 2 PCI Express x16 cards and as such drive the 4 outputs e.g. required by one 4K x 2K projector. Thus, LCD walls or projector arrays always require a graphics cluster composed of several high-end graphics PCs interconnected by a fast network like Infiniband or 10 Gigabit Ethernet. The simple task of displaying a single image on a tiled wall or on a projector array driven by a graphics cluster suddenly turns into a difficult parallel processing problem involving data replication and sophisticated synchronization.

This becomes even more of an issue if we move from displaying static content like gigapixel images to interactive applications on high-resolution displays i.e. 3D realtime rendering. Large screen projection of standard resolution computer graphics has been used for some time. However, the typical pixel size for such a setup is about 2-4 mm which is an order of magnitude larger than the pixel size of a standard resolution monitor. Thus, the images displayed on those low-resolution

Powerwalls could only be viewed from some distance, the large screen did not convey more information, and reading small text was almost impossible. With the upcoming tiled displays of about 100 megapixels we expect nearly monitor resolution across the full width of a large wall. This will allow for new types of collaborative applications benefiting from the high-resolution wide field of view.

However, providing interactive applications which can take advantage of such large displays connected to a graphics cluster is a non-trivial problem. One approach is to combine the distributed framebuffers into one large virtual desktop as it is done in Xinerama for Linux systems or by specialized device drivers in Microsoft Windows environments. The idea is to broadcast the draw requests from an application running on a single system to all cluster nodes driving a display. As long as applications support arbitrary window sizes, they can then render into windows which spread over many tiles eventually covering the entire screen space. For exploiting distributed 3D graphics hardware, mechanisms like library preloading have to be used in addition in order to intercept the stream of graphics commands and to reroute it to the render nodes of the cluster. While this approach enables a traditional graphics application developed for a single screen to run in stereo on a large tiled display, it obviously does not scale with problem size since the application remains locked on one master node. An alternative approach would be a truly parallel graphics application which distributes its entire processing load across the cluster computers; however, the development of such programs is still at its infancy. An additional problem which has to be dealt with in any of the described scenarios is the user interface of such applications. Obviously, mouse-based interaction and classical menu selection are not appropriate for steering such applications and new interaction paradigms have to be investigated to support user groups in front of a large screen.

The application area which initiated and still drives the development of high-resolution displays is the interactive visualization of large datasets as they arise in computational science, from high-resolution sensors like in medical imaging, or from digital data collections like transaction logs. Over the last 20 years visualization has grown into a scientific discipline of its own at the crossroads of computer graphics, data analysis, and scientific computing. Prominent fields in scientific visualization, which primarily deals with 3D data with spatial context, are volume rendering, flow visualization, and terrain rendering. All of these areas made significant progress with respect to the size of the datasets they can visualize at interactive rates by exploiting the processing capabilities provided by modern graphics chips. Over the last five years, graphics processing units (GPUs) have progressed from devices specialized on generating pixels from 3D triangles into programmable and highly parallel computing engines. Top-of-the-line graphics cards available for a few hundred Euros from the retail store now provide more transistor functions and numerical performance than the multi-core CPUs sitting next to them on the main board. However, the parallel streaming architecture and the specialized shading languages require the development of new algorithms for unleashing this power.

The research focus of the Visualization Institute of the University of Stuttgart is on the development of such GPU-based visualization techniques and we will present several examples ranging from medical volume visualization, to the visualization of molecular dynamics, and terrain rendering with atmospheres and non-linear perspective. These techniques form the basis for any interactive visualization application dealing with realistically sized datasets and they are the necessary building blocks for parallel visualization systems driving future gigapixel displays.