A Powerful Tiled Display System with only ONE PC

Chee-Kien, Wong
Nanyang Technological University
Singapore
cheekien@horixon.com

1 Introduction

Tiled display systems usually employ PC clusters or specialized image generation systems that have multiple display channels. The purpose is to provide a high-resolution display of some three-dimensional information. We have designed a powerful and very low-cost tiled display system, in which a PC is used with several commercial-of-the-shelf (COTS) graphics cards (nVidia's GeForce2MX) to produce the similar display effect as of those from expensive systems.

This novelty of this design lies in the parallel rendering software framework, which is capable of distributing rendering load amongst the graphics cards, thereby achieving up to three times greater polygon throughput compared to a conventional PC. Although operating systems like Windows support multiple graphics cards for extended display space, it inherently does not implement parallel processing by distributing the rendering load. Each graphics card in the latter case is simply allocated with the entire scene information.

2 Hardware

One definite requirement of tiled display systems is synchronized, multiple display channels. The most convenient way to achieve this, if not through dedicated image generation hardware, is to use networked PCs. This is because each PC's graphics card contributes a tile (or a channel) of the entire display wall. Our rendering system consists of four COTS PCI graphics cards in one PC. Our project has also addressed two important concerns arising from this setup: (i) the performance of PCI versus AGP graphics cards and (ii) how well the rendering performance scaled as we employed up to four graphics cards in the PC.

One advantage of our system is that there are no network communication overheads since all processing is done local to the PC. Four projectors were used in our tiled display system. They were connected to the display output of each graphics card in the PC.

3 Software

A parallel rendering software is designed to effectively make use of this hardware setup and it has dealt with the following two major issues in parallel rendering effectively: (i) load sorting and distribution and (ii) overheads from rendering of overlapping polygons.

The load distribution mechanism adopts the *Sort-First* strategy. In the *Sort-First* approach, graphics primitives are distributed amongst the rendering processors, which in this context, refers the different graphics cards in the PC. For brevity, *Sort-Middle* and *Sort-Last* algorithms cannot be implemented on the PC platform due to hardware constrains. The *Sort-First* approach is most suitable for our parallel rendering framework, since each COTS graphics card can execute the entire rendering pipeline independently after primitives are allocated to it.

A quad-tree screen-space partitioning algorithm is adopted to distribute the graphics load amongst the processors at each frame. This scheme is efficient because each processor is allocated only a quarter of the screen space and renders the primitives that fall within its boundaries. This quad-tree design also minimizes the overheads of rendering redundant overlapping polygons based on its hierarchical division of the screen space. Processors share the rendering load by communicating via simple message structures rather than passing primitives around.

A load manager distributes additional rendering jobs to "under-loaded" processors and the rendering of these jobs is done asynchronously with its main load by using threads. Completed rendering jobs are sent to the system memory and blitted back to the originating processor's framebuffer memory. These additional rendering jobs constitute less than a quarter of the screen space and thus are relatively cheap to be transported across the system bus among the graphics cards. To avoid tearing effects, all threads handling rendering processes are carefully synchronized by event-calls.

4 Results

We ran both synthetic and real-to-life datasets on this parallel rendering machine. The results show that the polygon throughput can be scaled up to three times higher (4xGF2MXs-18 MPoly/s, 2002) than a conventional PC that uses only one graphics card. This makes the system suitable for rendering high polygon-count datasets or 3D environments. An additional benefit of this system lies in its flexibility to be upgraded with the latest graphics hardware, thereby keeping the rendering power scalable.

By virtue of its four display outputs, this system also produces a display resolution of four times higher than the output of a conventional PC. We see the promising future of this system's implementation especially with the forthcoming PCI-Express technology, by which the graphics cards' combined performance may be further improved. Future work may be done to network such PCs to further scale the resolution and rendering power of the tiled display system.

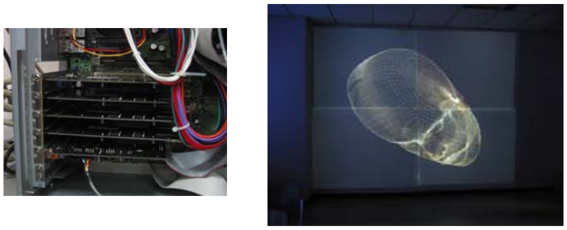

Figure 2: The PC setup and the actual tiled display wall (showing a high polygon-count skull model). Cost of system is approximately US$1500 only (excluding projectors).

References

Molnar, S., Cox, M., Ellsworth, D. and Fuchs, H., "A Sorting Classification of Parallel Rendering," IEEE Computer Graphics and Applications, Vol. 14, No. 4, pp 23-32, July 1994.