A Scalable Hardware and Software System for the Holographic Display of Interactive Graphics Applications

Tibor Balogh[1], Tamás Forgács[1], Tibor Agócs[1], Olivier Balet[2], Eric Bouvier[2], Fabio Bettio[3], Enrico Gobbetti[3] and Gianluigi Zanetti[3]

[1] Holografika, Budapest, Hungary – www.holografika.com – {t.balogh|t.forgacs|t.agocs}@holografika.com
[2] CS Communication & Systèmes, Paris, France – www.c-s.fr – {Olivier.Balet|Eric.Bouvier}@c-s.fr
[3] CRS4, Cagliari, Italy – www.crs4.it/vic/ {Fabio.Bettio|Enrico.Gobbetti|Gianluigi.Zanetti}@crs4.it

Abstract

We present a scalable holographic system design targeting multi-user interactive computer graphics applications. The display uses a specially arranged array of micro-displays and a holographic screen. Each point of the holographic screen emits light beams of different color and intensity to the various directions, in a controlled manner. The light beams are generated through a light modulation system arranged in a specific geometry and the holographic screen makes the necessary optical transformation to compose these beams into a perfectly continuous 3D view. With proper software control, the light beams leaving the various pixels can be made to propagate in multiple directions, as if they were emitted from physical objects at fixed spatial locations. The display is driven by DVI streams generated by multiple consumer level graphics boards and decoded in real-time by image processing units that feed the optical modules at high refresh rates. An OpenGL compliant library running on a client PC redefines the OpenGL behavior to multicast graphics commands to server PCs, where they are re-interpreted for implementing holographic rendering. The feasibility of the approach has been successfully evaluated with a working hardware and software 7.4M pixel prototype driven at 10-15Hz by three DVI streams.

Categories and Subject Descriptors (according to ACM CCS): B.4.2 [Input/Output and Data Communications]: Input/Output Devices Image Display

1. Introduction

In this short paper contribution, we briefly present a scalable holographic system design targeting multi-user interactive computer graphics applications. The display uses a specially arranged array of micro-displays and a holographic screen. Each point of the holographic screen emits light beams of different color and intensity to the various directions, in a controlled manner. The light beams are generated through a light modulation system arranged in a specific geometry and the holographic screen makes the necessary optical transformation to compose these beams into a perfectly continuous 3D view. With proper software control, the light beams leaving the various pixels can be made to propagate in multiple directions, as if they were emitted from physical objects at fixed spatial locations. The display is driven by DVI streams generated by multiple consumer level graphics boards and decoded in real-time by image processing units that feed the optical modules at high refresh rates. An OpenGL compliant library running on a client PC redefines the OpenGL behavior to multicast graphics commands to server PCs, where they are re-interpreted for implementing holographic rendering.

The proposed solution is able to provide all the depth cues and is truly multi-user within a reasonably large field of view. The developed prototype display is already capable to visualize 7.4M pixels at 10-15Hz while providing horizontal parallax with 0.8° angular resolution within a 50° field of view. Its design, based on parallel components, is fully scalable, and the OpenGL based parallel renderer, that masquerades as a compliant OpenGL library, makes it possible to quickly develop holographic applications and to run in 'holographic' mode legacy applications based on the OpenGL standard. As highlighted in section 2, while certain other technologies share some of these properties, they typically do not meet our system's capability in all of the areas.

2. Related work

Developing a scalable holographic system targeting multi-user interactive computer graphics applications is a large engineering effort, that requires advances in a number of technological areas. A full survey is beyond the scope of this short paper. In the following, we just provide a brief overview of competing 3D display technology for naked eye users.

Autostereoscopic displays implement left/right eye separation using various optical or lens rasters directly on top of LCD or plasma screens (Sharp [EWO*95], IBM, DTI, Samsung). This type of display imposes a single static viewing position. To overcome these limitations, manufacturers of stereoscopic displays are developing head/eye-tracking systems capable of following the viewer's head/eye movement [WEH*98, RS00, PPK00]. However, such a solution cannot support multiple viewers and introduces latency.

Multi-view displays show multiple 2D images to multiple zones in space. They support multiple simultaneous viewers, restricting them, however, to be within a limited viewing angle. Multi-view displays are often based on an optical mask or a lenticular lens array. The Cambridge multi-view display is a classic design in this area [DML*00]. Recently, Mitsubishi [MP04] demonstrated a prototype based on this technology and assembled with sixteen 1024x768 projectors and a lenticular screen. A number of manufacturers (Philips [vPF96], Sharp [WHJ*00], Opticality [RR05], Samsung, Stereographics, Zeiss) produce monitors based on variations of this technology. Lenticular state of the art displays typically use 8–10 images, i.e., directions, at the expense of resolution. A 3D stereo effect is obtained when left and right eyes see different but matching information. The small number of views produce, however, cross-talks and discontinuities upon viewer's motion [Dod96]. Our solution, instead, presents a continuous image to many viewers within a large workspace angle, due to the high number of view-dependent pixels that contribute to a single image.

Volumetric displays project light beams on a semi transparent or diffuse surface positioned or moved in space, that scatters/reflects incoming light [MMMR00, FDHN01]. By proper synchronization, it is possible to reconstruct 3D objects (SeeReal [RS00], Actuality, Felix, Deep Video Imaging). Portrayed objects appear however transparent, since the light spots addressed to points in space cannot be occluded by foreground voxels.

Pure holographic displays generate holographic patterns to reconstruct the light wavefront originating from the displayed object, using acousto-optic materials [SHLS*95], optically addressed spatial light modulators [SCC*00], or digital micromirror devices [HMG03]. Compared to stereoscopic and multi-view technologies, the main advantage of a hologram is in the quality of the 3D reconstruction. These systems are still confined to research laboratories, since the fundamental principle imposes limitations on realistically achievable image sizes, resolution, speckle, with consequent narrow fields of view, alongside enormous computing capacity required to reach acceptable refreshment rates for true interaction. In current prototypes, the display hardware is very large in relation to the size of the image (which is typically a few centimeters in each dimension).

3. The holographic display

Our display is based on projection technology and uses a specially arranged array of micro-display projectors and a holographic screen. The projectors are used to generate an array of pixels at controlled intensity and color onto the holographic screen. Each point of the holographic screen then transmits different colored light beams into different directions in front of the screen. Similarly to what happens with holograms, each point of the holographic screen thus emits light beams of different color and intensity to the various directions, but in a controlled manner. The display is thus capable of reproducing an appropriate light fields for a given displayed scene. The light beams that compose the light field are generated by optical modules arranged in a specific geometry. Each module contains a micro-display and special aspheric optics. A high-pressure discharge lamp illuminates all the displays, leading to a brightness comparable to normal CRT displays.

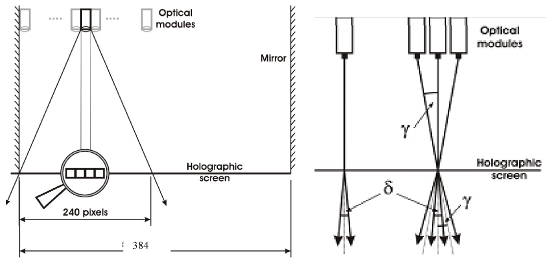

(a) Optical arrangement (b) Holographic screen

Figure 1: **Optical arrangement.** The light beams are generated through a modular light modulation system arranged in a specific geometry and the holographic screen makes the necessary optical transformation to compose these beams into a continuous 3D view.

The display system concept makes it possible to produce high pixel-count 3D images by optimizing the optical arrangement to the capabilities of the technology and the components applied. The prototype's overall 7.4M pixels originate from the resolution of the 96 LCD micro-displays, each of 320x240. The optical modules are densely arranged behind the holographic screen, and all of them project their specific image onto the holographic screen to build up the 3D image. Figure 1(a) illustrates the optical arrangement of the display.

The optical modules are not associated with specific view directions. Instead, the light beams to be emitted by the modules, i.e., the module images that are generated by the micro-displays, are determined by the geometry. Each module emits light beams toward a subset of the points of the holographic screen. At the same time, each point of the holographic screen is hit by more light beams arriving from different modules. In the current prototype, 96 optical modules project 240 pixels horizontally and 320 vertically. Each pixel on the screen is illuminated by 60 different LCDs, and the optical modules can be seen under different angles by looking from the pixel's point of view. This means that 60 different views are generated, and each view has 384x320 resolution. The imaging optics of the modules have a wide angle, which results in a 50° field-of-view. Since 60 independent light beams originate from each pixel in this field of view, the angular resolution of the display is 0.8°.

The holographic screen transforms the incident light beams into an asymmetrical pyramidal form. The cut of this light distribution is a long rectangle, where the vertical size of the rectangle is the vertical field of view, while the horizontal size corresponds to the neighboring emitting directions. This is the horizontal-only-parallax configuration. The principles on which the display is based would make it possible to provide vertical parallax. Doing so, would, however require another order of magnitude increase in data size, rendering times, and system complexity, for little gain in the visual performance in standard settings. We foresee exploring this option in the future, likely with different angular resolutions in the horizontal and vertical directions.

The horizontal light diffusion characteristic of the screen is the critical parameter influencing the angular resolution of the system, which is very precisely set in accordance with the system geometry. In that sense, it acts as a special asymmetrical diffuser. However, with standard diffusers and lenticulars it would be difficult to produce the shape of the required angular characteristics. The screen is a holographically recorded, randomized surface relief structure that enables high transmission efficiency and controlled angular distribution profile. These fully randomized (non-periodic) structures are non-wavelength dependent and eliminate moiré, without chromatic aberration. The precise surface relief structures provide controlled angular light divergence. The angular light distribution profile introduced by the holographic screen, with a wide plateau and steep Gaussian slopes precisely overlapped in a narrow region results in a highly-selective, low scatter hat-shaped diffuse characteristics. The prototype's holographic screen provides a horizontal angular diffusion of 0.8°, while the vertical diffusion is 60°. This means that the incident light beam's horizontal divergence will be 0.8°, and it is equal to the angle under which light beams are arriving from the neighboring modules. The result is a homogeneous light distribution and continuous 3D view with no visible crosstalk within the field of depth determined by the angular resolution. In figure 1(b), the horizontal light divergence of the incident light is δ and the angle under the neighboring optical modules is γ.

4. Parallel holographic rendering library

Interactive graphics applications are interfaced to the holographic display through a special implementation of OpenGL for holographic rendering. The library intercepts all OpenGL calls of the application. In addition to executing them on the local machine, using the native OpenGL library, it encodes each command into a command buffer and broadcasts it to the rendering back-end, which is responsible for holographic display. This is similar to cluster-parallel rendering in Chromium [HHN*02]. The rendering back-end consists in an array of PCs, connected to the display using one or more DVI connections. Each PC runs a server that controls an OpenGL framebuffer. The server is responsible for generating, starting from the original stream, the images associated to a fixed subset of the display rendering module. The server listens to the network and decodes the stream of multicast commands. Once decoded the commands are interpreted and transformed into native OpenGL commands and sent to the local OpenGL renderer. The interpretation of the commands involves modifying the way OpenGL command are generated according to parameters available from the local configuration service, to transform the original central view into the view associated with each of the associated optical modules. For each of the optical modules views, current frame's commands are re-executed, with the following modifications: the original perspective matrix is replaced with a matrix that matches the module's specific position and viewing frustum; a geometrical calibration is performed, to correct for non-linearities in the display/optical geometry; a light calibration is performed, to correct for the different color, contrast, and intensity response of the optical modules; an angular resolution correction (depth dependent anti-aliasing); intensity values belonging to directions in a given cone are averaged to match the screen diffusion characteristic and the actual light emission. The parameters required for each of these transformations are defined at configuration time.

5. Implementation and Results

We have implemented a prototype hardware and software system based on the design discussed in this paper. The developed small size prototype display is already capable of visualizing 7.4M pixels at 10-15Hz by composing optical module images generated by the 96 LCD displays. The displays are fed by image processing units, that decode the input DVI streams, realized on a Virtex II FPGA chip. The DVI channels work at 1280x1024 at 60Hz, and are thus capable of transmitting 225 MB/sec per channel. The current prototype uses three input channels for driving the display. The rendering library front-end runs on both Linux and Win32 operating systems, and currently implements most features of OpenGL 1.1. The library back-end, which drives the optical modules, is currently running on two Linux boxes equipped with GeForce6800GTS boards, that operate in twin-view mode. Three of the four outputs drive the display, while the fourth one is used for control purposes. Communication between front-end and back-end goes through a Gigabit Ethernet connection. The Linux boxes are connected to a Gigabit switch supporting IGMP snooping, so that the graphics command are efficiently multicast from the OpenGL client application. At present, back-end rendering time is the bottleneck, since each back-end PC has to render about fifty module images per frame. As a first step, we plan to exploit the recently introduced SLI GPU-teaming capabilities to boost back-end rendering performance.

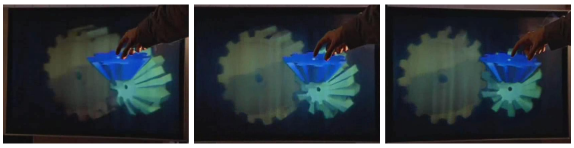

(a) OpenGL "gears" application

(b) Visualization of an abdominal aortic aneurysm reconstructed from CT data

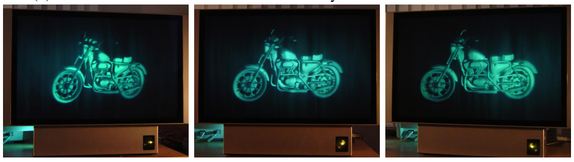

(c) Visualization of a geometric model

Figure 2: **Holographic display demonstration video: selected frames.** The video sequences were recorded using a hand held videocamera.

It is obviously impossible to fully convey the impression provided by the display on paper or video. An accompanying video show sequences of static and dynamic scenes recorded live using a moving camera. Representative video frames are included in figure 2. One of the examples is the OpenGL "gear" demonstration program, running without changes on the holographic display, thanks to dynamic linking with our OpenGL rendering library. All sequences were recorded with a hand held video camera. The operator was freely moving inside the display workspace. Note the parallax effects and the good registration between displayed object space and physical space, which demonstrate the multi-user capability of the display. As demonstrated by the video, the perceived image is fully continuous. This is qualitatively very different from other contemporary multi view display technology, that forces users into approximately fixed positions, because of the abrupt view-image changes that appear at the crossing of discrete viewing zones [Dod96]. By contrast, our display provides continuous horizontal parallax with 0.8° angular resolution for the full 50° field of view. This is about an order of magnitude better than current state-of-the-art multi view displays.

6. Conclusions and Future Work

We have presented a design and prototype implementation of a scalable holographic system design, that targets multi-user interactive computer graphics applications. The current display prototype is already sufficient for developing compelling prototype 3D applications that exploit its truly multi-user aspects. We are currently working on two demonstrators: one for the medical market (CT data analysis), and one for the CAD market (design review). These applications will be the driving forces for the design of our next generation display, currently under development, that targets the rendering of the equivalent of 50M pixels at interactive rates. It will be a large-scale 3D system with screen diagonal size of 1.8 meters, and a pixel reduced by 15% relative to the current models, to enable displaying high resolution 3D images.

Acknowledgments. This research is partially supported by the COHERENT project (EU-FP6-510166), funded under the European FP6/IST program.

References

[DML*00] DODGSON N. A., MOORE J. R., LANG S. R., MARTIN G., CANEPA P.: Time-sequential multi-projector autostereoscopic 3D display. *J. Soc. for Information Display 8*, 2 (2000), 169–176. 2

[Dod96] DODGSON N. A.: Analysis of the viewing zone of the cambridge autostereoscopic display. *Applied Optics: Optical Technology & Biomedical Optics 35*, 10 (1996), 1705–1710. 2, 4

[EWO*95] EZRA D., WOODGATE G. J., OMAR B. A., HOLLIMAN N. S., HARROLD J., SHAPIRO L. S.: New autostereoscopic display system. In *Stereoscopic Displays and Virtual Reality Systems II* (1995), vol. 2409 of *SPIE proceedings*, pp. 31–40. 2

[FDHN01] FAVALORA G., DORVAL R., HALL D., NAPOLI J.: Volumetric three-dimensional display system with rasterization hardware. In *Stereoscopic Displays and Virtual Reality Systems VII* (2001), vol. 4297 of *SPIE Proceedings*, pp. 227–235. 2

[HHN*02] HUMPHREYS G., HOUSTON M., NG R., FRANK F., AHERN S., KIRCHNER P., KLOSOWSKI K.: Chromium: A stream-processing framework for interactive rendering on clusters. In *SIGGRAPH 2002 Conference Proceedings* (2002), Hughes J., (Ed.), Annual Conference Series, ACM Press/ACM SIGGRAPH, pp. 693–702. 3

[HMG03] HUEBSCHMAN M., MUNJULURI B., GARNER H.: Dynamic holographic 3-d image projection. *Optics Ex-press 11* (2003), 437–445. 2

[MMMR00] MCKAY S., MAIR G., MASON S., REVIE K.: Membrane-mirrorbased autostereoscopic display for teleoperation and telepresence applications. In *Stereoscopic Displays and Virtual Reality Systems VII* (2000), vol. 3957 of *SPIE Proceedings*, pp. 198–207. 2

[MP04] MATUSIK W., PFISTER H.: 3D TV: a scalable system for real-time acquisition, transmission, and autostereoscopic display of dynamic scenes. *ACM Transactions on Graphics 23*, 3 (Aug. 2004), 814–824. 2

[PPK00] PERLIN K., PAXIA S., KOLLIN J. S.: An autostereoscopic display. In *Siggraph 2000, Computer Graphics Proceedings* (2000), Akeley K., (Ed.), Annual Conference Series, ACM Press / ACM SIGGRAPH / Addison Wesley Longman, pp. 319–326. 2

[RR05] RELKE I., RIEMANN B.: Three-dimensional multiview large projection system. In *Stereoscopic Displays and Virtual Reality Systems XII* (2005), vol. 5664 of *Proc. SPIE*. 2

[RS00] ROBERTS J. W., SLATTERY O.: Display characteristics and the impact on usability for stereo. In *Stereoscopic Displays and Virtual Reality Systems VII* (2000), vol. 3957 of *SPIE proceedings*, p. 128. 2

[SCC*00] STANLEY M., CONWAY P., COOMBER S., JONES J., SCAT-TERGOOD D., SLINGER C., BANNISTER B., BROWN C., CROSSLAND W., TRAVIS A.: A novel electro-optic modulator system for the production of dynamic images from giga-pixel computer generated holograms. In *Practical Holography XIV and Holographic Materials VI* (2000), vol. 3956 of *SPIE Proceedings*, pp. 13–22. 2

[SHLS*95] ST.-HILLAIRE P., LUCENTE M., SUTTER J., PAPPU R., SPARRELL C. G., BENTON S.: Scaling up the mit holographic video system. In *Proc. Fifth International Symposium on Display Holography* (1995), SPIE, pp. 374–380. 2

[vPF96] VAN BERKEL C., PARKER D., FRANKLIN A.: Multiview 3d-lcd. In *Stereoscopic Displays and Virtual Reality Systems III* (1996), vol. 2653 of *SPIE proceedings*, p. 32. 2

[WEH*98] WOODGATE G. J., EZRA D., HARROLD J., HOLLIMAN N. S., JONES G. R., MOSELEY R. R.: Autostereoscopic 3d display systems with observer tracking. *Image Communication - Special Issue on 3D Video Technology (EURASIP - 1998)* (1998), 131. 2

[WHJ*00] WOODGATE G. J., HARROLD J., JACOBS A. M. S., MOSELEY R. R., EZRA D.: Flat-panel autostereoscopic displays: characterisation and enhancement. In *Stereoscopic Displays and Virtual Reality Systems VII* (2000), vol. 3957 of *SPIE proceedings*, p. 153. 2