A Low-Cost Projection Based Virtual Reality Display

Dave Pape*[a], Josephine Anstey[b], Greg Dawe[a]

[a]University of Illinois at Chicago, Chicago, IL
[b]University at Buffalo, Buffalo, NY

ABSTRACT

This paper describes the construction of a single screen, projection-based VR display using commodity, or otherwise low-cost, components. The display is based on Linux PCs, and uses polarized stereo. Our aim is to create a system that is accessible to the many museums and schools that do not have large budgets for exploring new technology. In constructing this system we have been evaluating a number of options for the screens, projectors, and computer hardware.

Keywords: Projection based virtual reality, passive stereo, PC graphics

1. INTRODUCTION

Projection based virtual reality systems have become well established in application areas such as computational science, automotive engineering, and chemical exploration. These fields are often favored with large budgets and can afford expensive, advanced displays. VR also has applications for art, cultural heritage, and educational institutions, many of which have much smaller budgets, or are not able to support and maintain high-end graphics workstations. A simpler, more affordable projection based display system would be valuable for these institutions.

Currently, there are about 5 CAVEs and a handful of other virtual reality displays in public museums worldwide. Their popularity with the general public and the economics of throughput mean that relatively large groups are shown short demonstrations. Because of the limited time it is unusual for visitors to interact directly with the experience — very often a museum worker will navigate and interact for them. We believe that a cheaper system would allow visitors longer, more hands-on, and possibly more intimate experiences of virtual worlds.

This paper describes the construction of a single screen, passive stereo, VR display based on commodity, or otherwise low-cost, components. The display was constructed at the Department of Media Study, of the University at Buffalo, and is being used in teaching new media courses and for art exhibitions. We will identify the requirements that drove various design decisions, and compare the quality of the new display with that of high end systems like a CAVE or ImmersaDesk.

2. BACKGROUND

Many research groups have been working on producing PC-driven projection based VR displays. Some of these have focused on reducing the cost of such displays, while others seek to develop high-end systems using more widely available computers, to ease application development.

The NAVE was designed as a multi-screen display with many of the features of a CAVE, but at a lower cost.[1] It is a 3 screen, PC-driven, passive stereo display. It sacrifices the CAVE features of a floor projection and tracking. Four PCs are used to generate and synchronize the graphics, which are displayed using linear polarization via VREX 2210 projectors.

Belleman et al. describe the construction of a single-screen active stereo system.[2] Their Linux Immersive Environment uses a single Linux PC with a video sync-doubler to generate active stereo graphics. The other

* E-mail: dave.pape@acm.org

components — LCD shutter glasses and tracking — are equivalent to those in older single-screen projection based displays.

The AGAVE display was developed in parallel with the system described in this paper.[3] Its purpose is to augment Access Grid[4] displays to allow networked collaborators to share three dimensional content. It uses a front-projected, passive stereo display, and optionally includes tracking.

Z-A Productions market the SASCube as a PC-based CAVE-clone.[5] It is a full four screen, tracked, active stereo system driven by a cluster of PCs. The PCs use workstation class graphics cards (3DLabs Wildcats) in order to support active stereo, video genlocking and swap synchronization of the four walls. The intended markets for the SASCube include museums and public exhibitions.

The Ars Electronica Futurelab are developing the ArsBox PC CAVE.[6] The ArsBox is driven by Linux PCs using sync-doublers for active stereo video. Its purpose is to allow the Futurelab to more easily and affordably deploy art and technology VR applications that they develop for clients.

The HyPI-6 at the Fraunhofer IAO Virtual Reality Laboratory is a dual-platform, 6 sided CAVE-like system.[7] It can be driven by a large SGI Onyx computer, or by a cluster of PCs. The system can also function in both active stereo (when using the Onyx) and passive, polarized stereo (when using the PC cluster).

3. SYSTEM CONSTRUCTION

The core elements of the system are three PCs, two LCD projectors, tracking, and polarized stereo. The main computer, for the graphics, is a dual processor Linux PC with a two-channel 3D video card. The tracking system is a PC with an Ascension SpacePad, and a wand interactive device. The audio system is a PC with a generic audio card and speakers (and optionally a mixer and more expensive speakers). All the PCs are connected by Ethernet. The stereo display uses circularly polarizing filters for the two projectors and inexpensive polarized glasses. Figure 1 shows the system in use.

A full CAVE is an expensive device for several reasons. Typically the most obvious cost is the large, multi-pipe graphics computer required to drive it. Other major costs are the four (or more) CRT projectors and the CAVE structure itself. One significant, but often overlooked, cost is the architectural requirements. A 10' x 10' x 10' CAVE requires roughly 30' x 20' of floorspace for the wall projectors, and 14' or higher ceilings for the floor projector[8] (five and six wall CAVEs require even more vertical space). Obtaining this much space for a single VR device can be extremely problematic in many institutions, unless a very large amount of funding is involved. The ImmersaDesk and Responsive Workbench,[9] being only single screen displays, are more practical for many users because they require much less space. Similarly, we have restricted our system to just one screen.

More space could be saved by using front projection. However, with a front projected display, if viewers approach the screen too closely, they will cast shadows on it and obstruct their own view. Since we intend the system to be used for applications that feature direct, physical interaction with the virtual environment, we must still use rear projection.

3.1. Passive stereo

One of the most significant changes in this system, in comparison to an ImmersaDesk, is the use of passive, polarized stereo. To achieve the stereo effect, two projectors are used for the single screen, one for each eye's view. Differently polarizing filters are placed in front of each projector lens, and then users wear polarizing glasses where each lens only admits the light from the corresponding projector.

There are a few reasons for the switch to passive stereo. First, it allows us to use LCD projectors. The traditional active stereo approach requires high video frequencies — 96Hz or higher in order to guarantee that each eye sees a flicker-free image. This in turn requires the use of CRT projectors, as LCD and DLP projectors cannot handle these high frequencies (except for new, very expensive DLPs). Good CRT projectors are very expensive, whereas LCD projectors are very common and much more affordable. In addition, LCD projectors can be much brighter than CRT ones; one major flaw of the CAVE has always been its dimness.

Figure 1: The system in operation

In addition to the cheaper projectors, polarized stereo glasses cost much less than the LCD shutter glasses used in active stereo. The polarized glasses are also less fragile, which is important in a museum or other public environment, where they will be subject to rougher handling than in a research lab.

The polarization for passive stereo can either be linear (where the light for one eye is polarized at right angles to light for the other eye) or circular (where the light for one eye has left-hand polarization and that for the other has right-hand polarization). Linear polarizers are more common; however, when linear polarization is used, viewers must keep their heads level or they will lose the stereo effect as the polarization of their glasses no longer matches that of the projector filters. Because we want users to be able to interact freely with the system without affecting the stereo, we chose circular polarizing filters.

3.2. Screen material

Passive stereo in rear-projected systems has been difficult in the past because standard rear-projection screens do not preserve polarization; after evaluating numerous screen materials, we have identified one that produces very little stereo crosstalk (a.k.a. ghosting).

We obtained a number of different screen material samples from two manufacturers (Stewart Filmscreen and Da-Lite), and tested them with polarizing filters to determine how bright images would appear and how visible the stereo crosstalk would be. Table 1 summarizes the results of these tests. The tests were performed using a 35mm slide projector as directional light source 1 meter from the screen sample; a cropping slide was used to limit the illumination to a 2" x 3" area of each sample. A Gossen LunaPro light meter set to ASA 50 was used to take readings, with the diffuser in place, less than 1" from the screen's surface. The ambient light level reading was EV = -5 (less than 0.010 footcandles).

For the first column of readings, a left hand circular filter was placed at the projector lens. This is a general indicator of the relative brightness of each screen material. For the second reading, a second left hand circular filter was placed at the viewer's side of the screen sample. This is a relative measure of the amount of light reaching the eye through the passive glasses. The third reading is taken with a right hand circular filter replacing the left hand filter at the viewer's side of the screen sample. This is a relative measure of the amount of extinguished light reaching the eye through the passive glasses. The delta number is obtained by subtract reading 3 from reading 2. It is presented in f-stops. The larger the f-stop difference, the greater the extinguishing effect (less crosstalk).

The material we chose for our screen is referred to as "Disney black diffusion film" (Dbd). The Dbd material's performance is superior to all the other tested offerings, both because of its minimal stereo crosstalk and because of its contrast qualities. Specifically, the matte black surface of Dbd affords excellent ambient light absorption. This, coupled with brighter LCD projectors, allows acceptable image contrast with only subdued room illumination. For publicly attended showings, increasing the general room lighting has long been considered a desired safety improvement. Further, the material's neutral density effect on the transmitted light masks the "light bleed of black" typical of LCD based projectors. Image contrast is thereby further enhanced.

Table 1. Screen material test results using circular polarizing filters. Brightness (in f/stops and foot-candles) of a test projection at the screen, through the correct lens filter, and through the opposite lens filter.

Material	Screen	View	Ghost	Delta
Stewart "Disney" black [0.020" membrane]	5.0 [16fc]	4.5 [12fc]	-1.0 [0.26fc]	5.5 f/stops
Stewart '96 sample, (black) [0.020" membrane]	4.0 [8fc]	3.5 [6fc]	-1.5 [0.2fc]	5.0 f/stops
Stewart FilmScreen 200 regular (gray) [membrane on 3/8" plexiglass]	6.5 [48fc]	6.0 [32fc]	1.3 [1.3fc]	4.7 f/stops
Stewart FilmScreen 200 HiGain (gray) [membrane on 3/8" plexiglass]	6.5 [48fc]	6.0 [32fc]	1.0 [1fc]	5.0 f/stops
Stewart FilmScreen 150 (gray) [membrane on 3/8" plexiglass]	6.5 [48fc]	5.7 [28fc]	2.0 [2fc]	3.5 f/stops
DA-WAN-HC [gray coating on 3/8" black glass]	4.5 [12fc]	4.0 [8fc]	0.5 [0.35fc]	3.5 f/stops
DA-TEX [gray 0.011" membrane]	6.5 [48fc]	6.0 [32fc]	4.0 [8fc]	2.0 f/stops
DA-KS150 [gray coating on 3/8" black glass]	5.3 [20fc]	4.3 [10fc]	3.0 [4fc]	1.3 f/stops
DA-KS150 [gray coating on 3/8" plexiglass]	5.5 [24fc]	4.5 [12fc]	3.5 [6fc]	1.0 f/stops

3.3. Computer

The main computer for the system is a dual-processor Linux PC. We chose the Linux operating system because it is compatible with much of the existing CAVE VR software tools developed for Irix. In particular, SGI's OpenGL Performer[10] toolkit is the basis for most of our current applications, and is only available for Irix and Linux. The primary drawback of Linux, compared to Microsoft Windows operating systems, is limited hardware support; in particular, fewer 3D accelerated graphics cards are supported. However, in the case of the dual-channel graphics cards that are needed for this system (discussed more in the next section), almost all consumer level cards are currently supported under Linux.

We use a computer with two CPUs because of the importance of multi-processing for reliable real-time performance; the main computer must perform application calculations and render the graphics, as well as communicating with the tracking and audio PCs.

The specific computer we are using for our prototype system has two 700 MHz Pentium III CPUs, 256 MB of 100 MHz memory, and a Matrox G450 AGP graphics card.

3.4. Graphics cards

The passive stereo display requires two separate channels of video, one for the left-eye view and one for the right-eye view. As high performance 3D cards for PCs use the AGP interface, and no PC motherboards exist with more than one AGP port, to get two channels of video from a single PC requires a dual-channel graphics card. We preferred to stick with a single graphics computer in order to keep the system simple. Although it is possible to construct a PC cluster system to render multiple channels, the software tools for programming such a cluster were not in an advanced enough state for easy, general use (although this is rapidly changing). Keeping the system simple to both program and maintain is important for making it accessible to institutions other than VR research groups.

A handful of dual-channel PC graphics cards are currently available. At the time that we constructed the prototype system, we used the Matrox G450 because it was the only consumer level 3D card that had support for two channels with 3D acceleration under Linux (the Evans & Sutherland Tornado 3000 was another option, but it cost significantly more money without providing significantly greater performance). Since that time, other dual-channel cards have come out with Linux support, and we have tested them in other systems for comparison. Table 2 summarizes the computers and graphics cards tested.

Table 2: Specifications of the different computers tested

Graphics	Computer
Matrox G450	2 700 MHz Pentium IIIs, 256 MB memory (100 MHz), RedHat 7.1 Linux, Xi Graphics LGD X server 2.0 beta
nVidia Quadro2MXR	2 933 MHz Pentium IIIs, 256 MB memory (133 MHz), RedHat 7.0 Linux, XFree86 4.0.1, nVidia 0.9-769 OpenGL driver
nVidia GeForce2MX	2 550 MHz Pentium IIIs, 256 MB memory (100 MHz), RedHat 6.2 Linux, XFree86 4.0.2, nVidia 0.9-769 OpenGL driver
ATI Radeon VE	500 MHz Pentium III, 256 MB memory (100 MHz), SuSE 6.4 Linux, Xi Graphics LGD X server 2.0 beta
SGI InfiniteReality	Onyx2, 8 195 MHz R10000s, 2 GB memory, IRIX 6.5

Figure 2 and table 3 show the results of two sample benchmark tests. The primary benchmark used was SPECglperf 3.1.2.[11] SPECglperf is a toolkit that measures the performance of individual low-level OpenGL rendering primitives. It provides a way to characterize the relative performance of different graphics hardware over a range of possible drawing options. The plot in figure 2 shows the results of one specific SPECglperf test, that of drawing display-listed, Z-buffered triangle strips, with the size of the individual triangles ranging from 1 to 512 pixels. The results are representative of the results of most of the tests — although the PC cards are mostly slower than an Onyx2/IR, the better ones are comparable, and can even exceed the Onyx2 in simpler cases.

The results in table 3 are from an actual CAVE application adapted as a benchmark. The application used is Crayoland, a widely distributed program that makes heavy use of texture mapping on simple polygons to create a complex visual environment at interactive rates; it includes a few dozen autonomous objects moving under simple rules and several ambient sampled sounds. The table shows the average frame rate for a single, monoscopic rendering of the environment, at 1024x768 resolution. Frame rate measurements were taken at three fixed viewpoints for fifteen seconds each. The results are similar to those of the SPECglperf test; the Quadro2MXR card was able to outperform the Onyx2, while the Matrox and Radeon cards, although significantly slower than the Onyx2, were still able to perform at rates acceptable for interaction.

It must be noted that although the raw performance of current consumer PC graphics cards is quite good, the overall quality of the graphics is not always up to the level of an Onyx workstation. The Matrox G450 used in the prototype system lacks anti-aliasing support, and as a result images are sometimes very unpleasantly "jaggy". Newer cards include different forms of full-scene anti-aliasing, but they generally still do not have the resources to do this as well as an Onyx. Secondly, the best rendering performance on many of the PC cards is obtained by configuring them at a 16 bit color depth, rather than 24 or 32 bits. But, in this case they also use a 16 bit Z buffer, which can lead to severe rendering errors in scenes that feature both nearby and distant objects; eliminating these errors requires running the graphics card in its slower configuration.

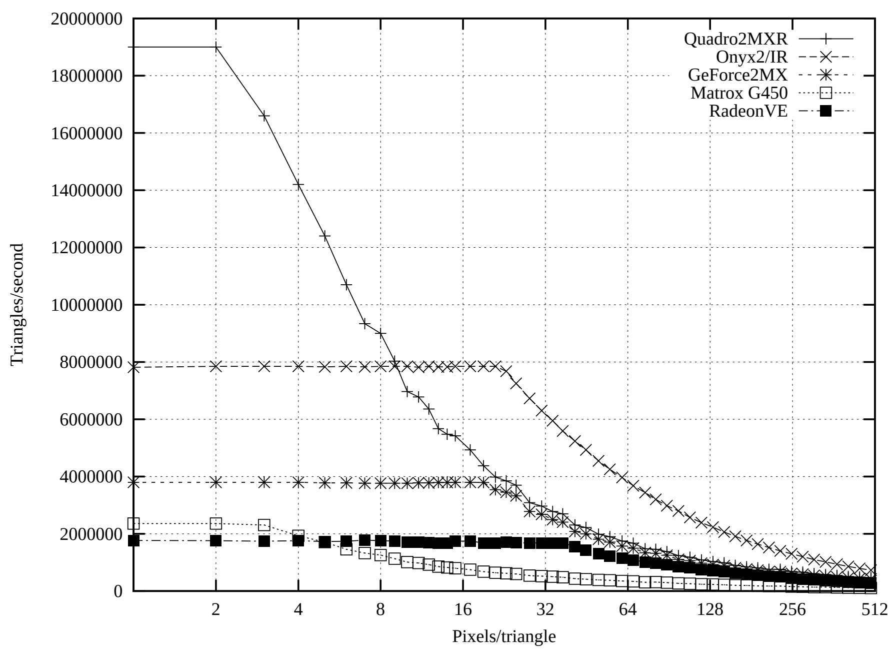

Figure 2. SPECglperf results for different graphics cards, drawing display-listed, Z-buffered triangle strips, with varying triangle sizes

Table 3: Performance of different PC graphics cards vs. SGI InfiniteReality, running a basic CAVE application

Hardware	**Average frame rate**
Matrox G450	31.0
Quadro2MXR	97.0
Radeon VE	32.6
Onyx2 InfiniteReality	78.6

3.5. Projector alignment

In order for the passive stereo to work, the images from the two separate projectors must match up on the screen. There are a few different approaches to this problem. Figure 3 illustrates four possibilities.

In our prototype system, we use projectors with integral shift lenses. These allow one to move the projected image several inches up or down, making it possible to match up the images from two stacked projectors. However, projectors with this feature are significantly more costly than those without it (by roughly $3000to5000$ per projector, at recent prices). The other arrangements in figure 3 are for cheaper projectors with fixed lenses.

The second approach is based on the fixed keystone offsets used in many projectors designed for meeting room environments, where they can be placed on a table, or mounted to the ceiling, and project onto a standard presentation screen. By mounting one projector on the ceiling, and the second at table height, the light cones will intersect at a certain distance, where the screen can be placed. The intersection point can be varied by changing the vertical distance between the two projectors.

The third approach is to simply throw away some image resolution. Two projectors can be stacked, and then only the area where their light cones overlap is used for the computer display. This means using the lower portion of the upper projector's image, and the upper portion of the lower projector's image; this involves using correspondingly offset window viewports in the graphics software configuration. With relatively slim projectors and a large screen, the lost image area can amount to less than 10% of the vertical resolution.

A final approach is to accept some slight distortion in the images. The two projectors can be stacked, but not exactly parallel to each other. By tilting one slightly, the two image areas on the screen will roughly overlap. The tilted projector's image will suffer some keystone distortion in this case, so the two images will not match exactly. But, some initial, non-rigorous tests suggest that the error is not significantly large, and may be acceptable to the average user.

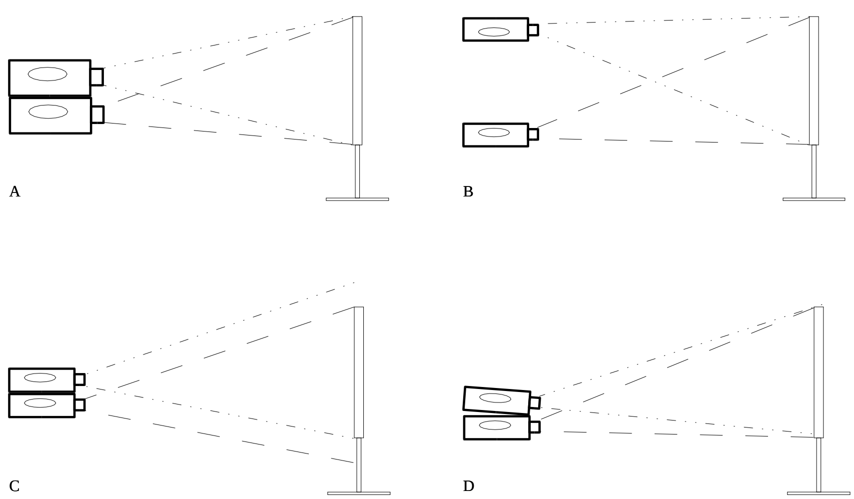

Figure 3. Projector configurations: A. Using shift lenses; B. Ceiling and table mounting; C. Throwing away some resolution; D. Tilting one projector

4. USING THE SYSTEM

This prototype system has been installed in the Department of Media Study, University at Buffalo. A group of graduate and undergraduate students helped build the system, and a current graduate project is to add spatialized sound capability to it. The department runs a two semester computer graphics course using OpenGL. In the second semester the students learn to create their own interactive, virtual environments which are displayed on the system. In a year long, more advanced, VR course for both graduates and undergraduates, teams of programmers and 3D modelers are building VR art projects. The students are developing their work using both the low-cost display in Media Studies and an ImmersaDesk at the Center for Computational Research across campus. Media Study has also established a lab with nine Linux PCs, loaded with the same software as that used on the VR system, on which students develop their projects. Because the VR system is PC-driven, students can easily run programs built on the other PCs, or even on their home computers, on the full VR display, rather than having to cope with the difficulties of moving between different platforms.

In the Spring of 2002, the system will be taken down and reassembled in Albany as one of the demonstrations at "UB day", a presentation by the University to the state legislature. In April 2002, we are planning to use the system (at Media Study) to show the students' VR projects at an event organized by Digital Poet, Dr. Loss Glazier. We have started preliminary talks with Hallwalls, a Buffalo media arts center, about staging a VR show using a similar system. We have also adapted the same technology (sans tracking) to create a monoscopic, two-screen display for an exhibit at the Smithsonian Museum's Arts and Industries Building. Working on these events is helping to clarify the problems involved when the system is exposed to the rigors of a museum or gallery environment, and to refine it in terms of durability and usability.

5. CONCLUSION

Affordable, PC-driven projection based virtual reality systems are a popular topic of investigation right now, and will probably soon become widespread. Our particular hope for such systems is that they will help expand VR out of the research and corporate labs, into public and educational venues.

Our prototype display has now been functional and in use for most of a year. The entire system cost roughly $20,000 to construct; we estimate that a new one could currently be built for about half that amount.

In basic performance tests, as well as day-to-day use, the low-cost PC system is comparable to one using an SGI Onyx2. The LCD projectors and black screen provide a bright display with better contrast than older systems using CRT projectors. The lightweight passive stereo glasses are less encumbering, and less fragile, than active glasses. The system as a whole can be maintained by a group of students who have only recently started learning about VR.

ACKNOWLEDGMENTS

We would like to thank Chris Galbraith, Dan Neveu and Paul Costa, for their enthusiasm and ingenuity in helping us put together the Media Study low-cost VR system.

The virtual reality research, collaborations, and outreach programs at the Electronic Visualization Laboratory (EVL) at the University of Illinois at Chicago are made possible by major funding from the National Science Foundation (NSF), awards EIA-9802090, EIA-9871058, ANI-9980480, and ANI-9730202, as well as the NSF Partnerships for Advanced Computational Infrastructure (PACI) cooperative agreement ACI-9619019 to the National Computational Science Alliance. EVL also receives major funding from the US Department of Energy (DOE) Science Grid program, awards 99ER25388 and 99ER25405, and the DOE ASCI VIEWS program, award B347714. In addition, EVL receives funding from Pacific Interface on behalf of NTT Optical Network Systems Laboratory in Japan.

REFERENCES

1. J. Pair, C. Jensen, J. Flores, J. Wilson, L. Hodges, and D. Gotz, "The nave: Design and implementation of non-expensive immersive virtual environment," in *Siggraph 2000 Sketches and Applications*, p. 238, 2000.
2. R. Belleman, B. Stolk, and R. de Vries, "Immersive virtual reality on commodity hardware," in *ASCI 2001 Conference*, 2001.
3. J. Leigh, G. Dawe, J. Talandis, E. He, S. Venkataraman, J. Ge, D. Sandin, and T. DeFanti, "Agave: Access grid augmented virtual environment," in *Proceedings of the Access Grid Retreat*, 2001.
4. Argonne National Laboratory / NCSA Alliance, "Access grid web page," (URL: http://www-fp.mcs.anl.gov/fl/accessgrid/).
5. Z-A Productions, "Sas cube web page," (URL: http://www.z-a.net/sascube/index.en.html).
6. Ars Electronica Future Lab, "Arsbox (pc cave) web page," (URL: http://www.aec.at/futurelab/homepage/show_pro.asp?PID=305).
7. Fraunhofer IAO Virtual Reality Lab, "Hypi-6 web page," (URL: http://vr.iao.fhg.de/6-Side-Cave/index.en.html).
8. M. Czernuszenko, D. Pape, D. Sandin, T. DeFanti, G. Dawe, and M. Brown, "The immersadesk and infinity wall projection-based virtual reality displays," in *Computer Graphics, Vol. 31 No. 2*, pp. 46–49, 1997.
9. W. Krueger and B. Froehlich, "The responsive workbench," in *Computer Graphics and Applications, Vol. 14 No. 32*, pp. 12–15, 1994.
10. J. Rohlf and J. Helman, "Iris performer: A high performance multiprocessing toolkit for real-time 3d graphics," in *Proceedings of SIGGRAPH '94 Computer Graphics Conference*, pp. 381–395, 1994.
11. SPECopc, "Specglperf web page," (URL: http://www.specbench.org/gpc/opc.static/glperf.htm).