CANVAS
A VIRTUAL REALITY ENVIRONMENT FOR MUSEUMS

Hank Kaczmarski
Beckman Institute
University of Illinois
Urbana-Champaign, Illinois USA
hank@isl.uiuc.edu

Kathleen Harleman
Krannert Art Museum
University of Illinois
Urbana-Champaign, Illinois USA
harleman@uiuc.edu

Abstract

***Canvas*, the Collaborative Advanced Navigation Virtual Art Studio is a scalable, reconfigurable display technology for modern art museums, not as a work unto itself, but as an environment to facilitate the creation and display of narrative art works. *Canvas* is collaborative because it can be connected to an array of geographically dispersed immersive virtual spaces, has advanced navigation to allow viewers in different locations to interact with virtual art and allows for the creation and presentation of virtual art, art that exists not in two or three-dimensional space like a painting or sculpture, but in the multi-dimensional world of virtual images.**

Immersive Art

The origins of room-sized immersive virtual reality art can be traced back to at least the mid-1990's at the Electronic Visualization Laboratory (EVL) on the campus of the University of Illinois, Chicago, Illinois, USA [1]. Building on EVL's invention of the CAVE™ (Cave Automatic Virtual Environment) in 1992 [2], early immersive art relied on the same powerful graphics supercomputers used in scientific visualizations, machines beyond the financial wherewithal of the vast majority of the art community and necessitated the co-location of the artist fortunate enough to have access to those supercomputers with the actual supercomputer. This co-location was doubly necessary since CAVE ™ No.1 was the only location in which to immerse the artist and audience in this newly created experience. For the next decade, the plight of the electronic artist was not significantly improved by technological advances. Multiple single-graphics-channel supercomputers were replaced by monolithic multiple-graphics-channel supercomputers paid for by government funding to advance the burgeoning field of scientific visualization. This only exacerbated the problem of access by artists to what had rapidly become national scientific treasures.

Into this environment of partially-immersive, supercomputer-driven, Unix-intensive immersive visualization, the Integrated Systems Laboratory (ISL) at the Beckman Institute of the University of Illinois, Urbana-Champaign, Illinois USA was formed [3]. Our initial mission was to support research by multi-disciplinary research teams of Perceptual and Cognitive Psychologists, Electrical and Computer Engineers, Computer Scientists, Mathematicians, Kinesiologists and Educators by creating and maintaining discrete projection-based environments, both immersive and high resolution. Toward this goal, the US National Science Foundation funded the construction of the world's most accurate fully-immersive virtual reality test room, the Beckman V-R Cube, constructed by TAN Projektionstechnologie of Dusseldorf, Germany, a six-surface active stereo environment powered initially by a Silicon Graphics twelve-processor three-InfiniteReality™-graphics-pipe supercomputer [4].

Once the dichotomy between the technology-resource-intensive sciences and the resource-hungry community of artists within the UIUC College of Fine and Applied Arts became too great to ignore, a Chancellor's initiative titled "Art in a Technology Intensive World" was started to allow access by artists to the same technology resources, including visualization facilities, formerly the exclusive venue of scientists. This now recreated the model started ten years earlier of sharing existing facilities among disciplines where access prioritization would be necessary due to the unique nature of these highly specialized, costly environments. It rapidly became obvious that this was not a sustainable model.

A research laboratory designed for psychological testing of human subjects is not an appropriate setting for creating and viewing art, even technologically intensive art. That only a gallery designed for electronic art would work for creation and presentation of electronic art was recognized by the authors of this paper. We set about to create such a gallery knowing the challenges were great in adapting technology developed by computer scientists to artists less proficient at programming in basic computer languages and previously funded at levels of national science priorities, levels not obtainable by the art community.

Goal:

Our goal was to create, within the confines of a traditional museum gallery space, an affordable, scalable, artist-as-programmer-friendly immersive projection environment, physically a room at least three meters on a side whose walls are constructed of rear-projection screen material. In its truest form, room-sized virtual reality is a one-person experience, since the stereo video images projected on the screen surfaces are drawn by the computers from one person's head and eyes perspective.

Stereo video, video with perceived depth, can be presented to a viewer's eyes by several methods. Active stereo, stereo that uses some device to switch images on and off at the viewer's eye or passive stereo which uses filters, either light-polarizing or color-selective, are two traditional approaches used to present differing images to each eye continuously or at least continuously enough to allow the viewer to experience visual depth. Active stereo (eyewear has an active component), at the room-sized implementation, requires a video projector capable of processing more than 96 complete video frames per second without dropping or interpolating any frame in the sequence, a very expensive projector. This frame-sequential (left eye image/right eye image) stream of images is decoded by electronic shutter glasses triggered by a graphics-computer-controlled array of infrared emitters, a very expensive, very delicate technology not terribly well suited to a minimally supervised, sustainable exhibit space. Active stereo is the choice of most current scientific immersive visualization work, but has financial implications and the eyewear is not well suited to people who have difficulty with stroboscopic imagery. Passive stereo (eyewear contains differing filters, but no electronically active element) relies on fairly inexpensive projectors (however twice as many as active stereo), typically presentation-grade projectors capable of 1024x768 pixel resolution with 2000 lumen light output. *Canvas* uses circularly-polarized passive stereo video technology, a method that is scalable to orthogonal surfaces, is faithful in color rendition, easy to implement and maintain, and relatively inexpensive.

This virtual reality space has, in a three-walled configuration, six projectors (two overlapping behind each "wall" of projection screen) to create images that appear to have depth, to leave the screen and occupy the space in front of or behind the screen. This is accomplished by sending different information to each of the viewer's eyes forcing the brain to see left and right eye images that are created to be slightly offset from one another, just as eyes, because they are offset in the horizontal plane, see features in real life. This imparts the characteristic of depth to the combined images, an effect called stereopsis, one of the strongest visual cues our primary visual cortex uses to add a third dimension to the planar viewing surfaces of *Canvas*. Scalable from single-surface displays to a completely immersive space using six or more surfaces, *Canvas* adapts to the scale and complexity of the art being presented.

Creating an easily reconfigurable screen/projection system requiring a minimum of change-over time while maintaining the integrity of the circularly polarized images presents mechanical and material-science issues. Most screen materials used for audio-visual presentations have "gain" which, while focusing light to give the audience a perceived brighter image, destroys polarization, as is also the case with most mirror systems used to reduce projector throw distances. As we wanted the participant to navigate within an immersive space, front projection of images, while conservative of overall space, limits the participant's usable space to often a very small area where the image would not be blocked by the participant's body. We then require all of the projectors to be behind the screens and the screen material itself to be able to pass light without distorting its polarization. Additionally, tradeoffs in adjacent surface image bleed against maximum light transmission limit the choices of screen material to just a few remaining screens, most notably the "Disney Black" series made by Stewart Filmscreen [5]. Most multi-chip LCD AV projectors polarize light internally at the combining prism, usually linearly with red/blue having the opposite polarization of green. This requires careful matching of projector and filter system, since even circularly polarizing filters often use a linear polarizer ahead of a ¼ wave retarder. Fortunately, mounting the circular polarizing filter set at a 45 degree angle to the horizontal mounting of the projectors will cancel much of this linear polarization, however at the cost of some additional light loss beyond the average 67% light loss of the filter alone. DLP projectors do not typically contain internal linear polarizing prisms, but have been observed to create moiré patterns on the light passing through the circular polarizing filters. Caution is needed in any case since different manufacturers' filters behave differently, so test the combination before purchasing. We recommend a combination of Epson 703/730 series of LCD projector with polarizing filters manufactured by American Polarizers.

Devices used to track the motion or position of audience members and haptic and spatialized sound feedback to audience members adds a significant sense of an individual experience to interactive or narrative art but syncronizing and range-limiting such devices in a multi-computer environment has been little explored before *Canvas.* Room-sized human motion tracking systems used to draw images and create spatial sounds based on a person's head and hand positions exist at reasonable prices for scientific visualization, but those costs and the fragile nature of some technologies caused us to consider alternatives for *Canvas*. Electromagnetic systems, such as those manufactured by Ascension Technologies [6] and Polhemus [7] require a transmitter sending multiple-orientation AC or pulsed DC signals to achieve six-degree-of-freedom (6DOF) information. Getting the 6DOF information into the graphics computers requires receivers on the head and hand of the viewer and cables or wireless transmitters, making the system cumbersome and expensive. Ideally, optical tracking would decode the head and hand positions automatically and advances are being made in that direction, but optical systems currently commercially available require retroreflective or active light emitting diode markers and numerous cameras, again an expensive solution. The current iteration of Canvas's software can accept and process this electromagnetic or optical information, but the fielded version of *Canvas* relies on compromises to the purist's view of a true immersive space in that the head position of the viewer is assumed to be in the center of the room volume and does not change as the viewer moves throughout the space. This allows multiple viewers to see a stable image scene unaffected by the "tracked" viewer, at the expense of not allowing a viewer to accurately explore the virtual objects based on head position and orientation. The problem of navigating through the virtual space is now able to be handled more easily and inexpensively with off-the-shelf technology. A wired or wireless gamepad connected through the USB port of a computer has sufficient digital and analogue controls to allow the viewer to navigate and change elements of the scene using a most likely already familiar input device.

Even using expensive active stereo CRT projectors, the greatest expense of a scientific immersive visualization space previously was the graphics computer(s) necessary to drive the projectors. Software to gather and process the tracking information and the availability of multiple cpu's with reasonable clock speed, massive centralized memory and stereo-capable graphics engines was available on specialized platforms from very few manufacturers, most notably Silicon Graphics. There was no real entry point for the artist community in immersive visualization until commodity personal computer manufacturers, spurred on by the computer video game industry, created a graphics card sector to rival the graphics engines of SGI. A cluster of personal computers, running any of the leading operating systems including Windows, OS X and Linux could easily demonstrate graphics benchmarks to interest a serious visual artist. A layer of software was still necessary above the operating system and below the application level to provide synchronization of input devices and control of the image buffers throughout the now-extremely-powerful independent graphics engines, each in its own computer. VRCO Incorporated, the current licensor of the software used to control the earliest Caves™, CaveLib™ and trackd™, licenses software for synchronization and control of clusters, but in limited operating systems, notably not OS X. It is possible to have even low level software overtake the cost of hardware in cluster environments so care needs to be taken at the initial stage of environment development to build a legacy of applications on software packages that minimize the expense of expansion within the virtual gallery space, that will allow for the electronic interconnection of one space with another at no additional software cost and that will allow artists to develop on their personal machines away from the gallery using the same software with no additional licensing fees. For museums that wish to provide an open development environment to artists while still maintaining intellectual property rights to the resultant electronic art, as much of the software as possible used in the display computers should be open source and ideally operating system agnostic. Universities have been leaders in scientific visualization from its inception and have maintained that lead by providing open source solutions to the challenge of controlling visualization clusters. Chromium [8], Net Juggler [9] and Syzygy [10] are three excellent examples of open source middleware designed to work in computer cluster environments. The images occupying the space of *Canvas* are generated in a cluster of personal computers, one computer having two video output channels per display surface, the two channels necessary to feed a left eye and a right eye projector. Each computer is custom-built using commodity components to assure upgradability yet is sufficiently minimalist to avoid cpu-clock-cycle-stealing "features". We particularly like a micro-ATX form factor motherboard with onboard gigabit network interface card (NIC) and fast AGP graphics slot. The cluster concept really works only because of the graphics engine, so selection of one is crucial. Several manufacturers of graphics cards having two outputs have software to combine those outputs into a single desktop which simplifies programming, as half of the screen width is dedicated to each eye and ultimately split in such a way to send each half to a separate output and hence to a separate projector. When multiple computers each have a graphics channel dedicated to creating images for one screen surface, it is necessary to coordinate the timing of the display of new images as they are created within the graphics cards. Each computer renders a new image based on the input data and requirements of the program being run, that image being stored in an off-line image buffer. When the last computer finishes drawing its corresponding image, software must tell all computers to switch all images simultaneously, forcing the cluster to perform as one graphics supercomputer. This synchronized buffer switching is called swaplocking, a necessary but much less rigorous requirement than the combination of buffer synchronization and scan locking (genlocking) necessary for active stereo systems. The ability of the motherboard to support (or have onboard) a second NIC is necessary for networked (geographically distributed) applications, one NIC used to locally communicate with other computers in the cluster and the second to communicate over the wide area network. For a multi-surface display, the resultant cluster of computers used in *Canvas* works together through Syzygy, software developed at the Integrated Systems Lab (ISL) at the Beckman Institute. This software supervises the actual application from the initial input data gathering responsibilities all the way through synchronization of the multiple computers' graphics channels swaplocking. Syzygy is freely available, supported by the ISL, and operates on Windows, Linux, Irix and OS X.

Initial Application

Canvas was adapted from scientific visualization and sonification hardware and software in stable operation at the ISL to support a narrative art piece *360*, a three-walled, gallery-centered immersive room presentation of the life-stages of the romantic life of a couple. In its initial showing at the Krannert Art Museum in August 2004, *360* contained a timeline of the couple meeting in a café, with the gallery viewer wearing stereo glasses seated at a real table within eye and earshot of the couple at a virtual table nearby. Navigation through the life of the couple was afforded through a joystick. The seated position of the viewer played off of the decision to have the three graphics computers draw all images and spatialized sounds from a fixed head position. Response from other artists and gallery attendees was so positive that we have made this system a permanent gallery installation, the foundation of an electronic gallery available to all artists.

Artist Entry

A computer programmer is a lousy paintbrush. We now have an environment in which virtual reality art can be created and displayed. But what knowledge of the programming languages underlying this gallery must an artist have and is it at all possible to insert an individual in the creative process to assist the non-programmer artist? *Canvas* is currently exploring these issues, because the environment allows several entry points into the ultimate art creation experience. Those proficient in computer programming in C/C++ or the Application Programming Interface OpenGL can write programs on their favorite operating system machines and re-compile on the console running *Canvas* or collaborate with a computer programmer possessing those skills. If the artist is knowledgeable in a computer scripting language, for instance Python [11], or has created works in a three-dimensional file format (VRML [12] for example), that level of entry is also supported. For creation of three-dimensional artwork with no programming requirements placed upon the artist, *Canvas* offers a full volume drawing program operated from a wireless wand (virtual paintbrush/chisel). Our efforts at advancing any or all of these entry modalities will depend on the level of artist acceptance of the ease of use versus programming power trade-offs inherent in each method.

References

[1] M. Dolinsky, “Creating Art Through Virtual Environments”, *Computer Graphics*, vol. 31, no. 4, New York, NY 11/01/1997 - 11/01/1997, ACM Press, pp. 34-35

[2] C. Cruz-Neira, Sandin, D., DeFanti, T., Kenyon, R., Hart, J., “The CAVE: Audio Visual Experience Automatic Virtual Environment”, *Communications of the ACM*, vol. 35, no. 6 06/01/1992, pp. 65-72

[3] http://www.isl.uiuc.edu

[4] http://www.isl.uiuc.edu/Labs/room_b650.htm

[5] http://www.stewartfilm.com/index2.htm

[6] http://www.ascension-tech.com/products/flockofbirds.php

[7] http://www.polhemus.com/fastrak.htm

[8] http://chromium.sourceforge.net/doc/LLNLcopy.html

[9] http://netjuggler.sourceforge.net/

[10] http://www.isl.uiuc.edu/syzygy.htm

[11] http://www.python.org/

[12] http://www.graphcomp.com/info/specs/sgi/vrml/spec/