

Technologies for Virtual Reality/Tele-Immersion Applications: Issues of Research in Image Display and Global Networking

Tom DeFanti, Dan Sandin, Maxine Brown,
Dave Pape, Josephine Anstey, Mike Bogucki, Greg Dawe, Andy Johnson
Electronic Visualization Laboratory (EVL)
University of Illinois at Chicago
851 S. Morgan St., Room 1120
Chicago, IL 60607-7053 USA
tom@uic.edu

Thomas S. Huang
Beckman Institute for Advanced Science and Technology
University of Illinois at Urbana-Champaign
405 N. Mathew Ave.
Urbana, IL 61801 USA
huang@ifp.uiuc.edu

Table of Contents

1.	Abstract	1
2.	Issues	2
2.A.	Background: Projection-Based VR Technologies	2
2.B.	Issues in Tele-Immersion Development	3
2.C.	Challenges of Tele-Immersion	4
3.	Correspondences/Dependencies: Tele-Immersive Device Design Concepts	7
3.A.	Motivation for Desktop/Office-Sized VR Display Devices	7
3.B.	New Immersive Display Technologies	7
3.C.	ImmersaDesk3	8
3.D.	Personal Augmented Reality Immersive System (PARIS)	9
3.E.	CyberCeilings, designed for the Last Unused Projection Surface	10
3.F.	Personal Penta Panel (P3) or Dilbert's Dream	11
3.G.	Totally Active Work Space (TAWS)	11
3.H.	CAVEscope: Simulating Variable Resolution Displays	12
4.	Questions for the Future	13
5.	Acknowledgments	13
6.	References Cited	14
7.	Biography	15

1. Abstract

The Electronic Visualization Laboratory (EVL) at the University of Illinois at Chicago (UIC) has developed an aggressive program over the past decade to partner with scores of computational scientists and engineers all over the world. The focus of this effort has been to create visualization and virtual reality (VR) devices and applications for collaborative exploration of scientific and engineering data. Since 1995, our research and development activities have incorporated emerging high bandwidth networks like the vBNS and its international connection point STAR TAP, in an effort now called *tele-immersion*.


As a result of eight years' experience building first and second-generation projection-based VR devices to support these applications, we wish to describe needed research in *third-generation* VR devices aimed at desktop/office-sized displays. Since no current projection technology is yet configurable with ideal resolution and size, we must first describe the variety of emerging display devices, such as large color plasma displays, LCD projectors, LED panels, Digital Light Valves (DLVs), Grating Light Valves (GLVs), and Digital Micro Mirror Displays (DMDs).

In 1991, we conceived and over several years developed the CAVE virtual reality theater, a room-sized, high-resolution, projection-based system that enables users to experience excellent immersion in full 3D imagery. We then developed the ImmersaDesk, a smaller, software-compatible, drafting-table-format version of the CAVE that has been deployed to dozens of locations, nationally and internationally, at government institutions, national laboratories, universities, and companies.

The hardware now needs to be made smaller, higher resolution and more adaptable to the human and his/her workspace. Middleware that manages connections, bandwidth and latency needs to be integrated with the computer systems driving these hardware devices. Software that increases the quality of human-computer interaction through human output recognition must be brought from specialized lab experiments to routine use, and provided as part of the tele-immersive collaborative experience. This paper discusses many of the issues at the heart of this research.

2. Issues

2.A. Background: Projection-Based VR Technologies



The CAVE™ is a multi-person, room-sized, high-resolution, 3D video and audio environment. Graphics are projected in stereo onto three walls and the floor, and viewed with stereo glasses. As a viewer wearing a location sensor moves within its display boundaries, the correct perspective and stereo projections of the environment are constantly updated, so the image moves with and surrounds the viewer to achieve immersion.



The ImmersaDesk™ is a drafting-table format version of the CAVE. When folded up, it fits through a standard institutional door, and deploys into a 6' x 8' footprint. It requires a single graphics engine of the SGI Onyx or Octane class, one projector, and no architectural modifications to the working space. The ImmersaDesk is software-compatible with the CAVE library.



The Infinity Wall is derivative of the PowerWall, a research effort of Paul Woodward at the University of Minnesota. The PowerWall achieves very high display resolution through parallelism, building up a single image from an array of display panels projected from the rear onto a single screen. High-speed playback of previously rendered images is possible by attaching extremely fast disk subsystems, accessed in parallel, to an Onyx. The Infinity Wall is a simpler PowerWall that has tracking and stereo; it is CAVE library compatible.

Computational Science and Engineering Research Partners. Since 1986, EVL has partnered with the National Center for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign and the Mathematics and Computer Science Division at Argonne National Laboratory in ongoing efforts to develop national collaborations at professional conferences—notably ACM SIGGRAPH and ACM/IEEE Supercomputing. These events emphasize high performance computing and communications, VR, and scientific visualization. The overall purpose is to encourage the development of teams, tools, hardware, system software, and human interface models on an accelerated schedule to enable national-scale, multi-site collaborations applied to National Challenge and Grand Challenge problems. As a result of I-WAY, an experimental high-performance network linking dozens of the USA's fastest computers and advanced visualization environments at Supercomputing 95, many successful CAVE collaborations resulted. [4, 5, 16] Dozens of these scientists continue to work with EVL on various joint research projects, either informally, through grants, or through affiliation in the NSF Partnerships for Advanced Computational Infrastructure (PACI) program. [<http://alliance.ncsa.uiuc.edu>] A number have CAVEs, ImmersaDesks and similar devices. The CAVE Research Network Users' Society (CAVERNUS) has been in operation for several years, and welcomes members interested in projection-based VR. [www.ncsa.uiuc.edu/VR/cavernus]

International STAR TAP Partners. STAR TAP is a persistent infrastructure to facilitate the long-term interconnection and interoperability of advanced international networking in support of applications, performance measuring, and technology evaluations. [www.startap.net] STAR TAP, in Chicago, is the Next Generation Internet Exchange (NGIX) point for Next Generation Internet (NGI) [www.ngi.gov] and Internet2 [www.internet2.edu] networks. Several institutions in foreign countries own CAVEs and ImmersaDesks or similar devices and have either obtained

STAR TAP connectivity or are applying for it in order to test broadband VR collaborations. Countries currently connected to STAR TAP include Australia, Canada, Japan, Korea, Taiwan, Singapore, Russia, Norway, Sweden, Denmark, Iceland, Finland, The Netherlands, France, and Israel. CERN is also connected.

Tele-immersion. The term *tele-immersion* was first used in October 1996 as the title of a workshop organized by EVL and sponsored by Advanced Network & Services, Inc. to bring together researchers in distributed computing, collaboration, VR, and networking. At this workshop, we paid specific attention to the future needs of applications in the sciences, engineering, and education. We define tele-immersion as the union of networked VR and video in the context of significant computing and data mining. EVL's Web site [www.evl.uic.edu] has an extensive tele-immersion bibliography and papers. Tele-immersion has since entered the NGI and Internet2 vocabulary. In the applications section of the Computing Research Association's "Research Challenges for the NGI," *tele-immersion* was one of five key technologies identified as necessary for the future use of the NGI [18]:

***Tele-immersion.** Tele-immersion will enable users in different locations to collaborate in a shared, virtual, or simulated environment as if they are in the same room. It is the ultimate synthesis of networking and media technologies to enhance collaborative environments. Tele-Immersive applications must combine audio, video, virtual worlds, simulations, and many other complex technologies. They will require huge bandwidth, very fast responses, and guarantees of delivery.*

We have connected CAVEs and ImmersaDesks over networks, from ATM-based 622 Mb and 155 Mb networks to ISDN. We have implemented video and audio over the networks to enable users to conduct remote teleconferencing and distributed virtual prototyping. At Supercomputing '97, we held a 17-way ImmersaDesk/CAVE tele-immersion experiment with 8 ImmersaDesks on the conference exhibit floor and another 9 devices connected from as far away as Amsterdam and Tokyo. [7] At Supercomputing '98, 10 countries participated in the iGrid booth showing many instances of international tele-immersion. [23, 24] [www.startup.net/igrd]

CAVERN is our acronym for the CAVE Research Network. CAVERN is comprised of dozens of network-connected CAVEs, ImmersaDesks, and other VR devices, like Head-Mounted Displays, Responsive Workbenches, and BOOMS. CAVERN is managed by the CAVE libraries and CAVERNsoft, a distributed shared memory software package optimized for networked collaboration. [12, 13, 14]

2.B. Issues in Tele-Immersion Development

The ideal tele-immersion system is not hard to imagine. Combine the best computer graphics, audio, computer simulation, and imaging. Connect with networking as good as direct memory access. Provide software and hardware to track gaze, gesture, facial expression, and body position. Offer it as a built-in feature on all personal computers and video games. Obviously, we are far from achieving ubiquitous tele-immersion.

Consider human voice and audio in general. There is a worldwide network optimized for speech (the telephone system) that supports 2-way and multi-way interactions. Computers and other equipment one can purchase in shopping malls can record, edit, playback, and duplicate (even net broadcast) audio to perfection. Real-time speech synthesis is close at hand with gigaflop desktop machines. Similarly, for standard video, recording, editing, playback, global teleconferencing, and broadcast, mature and optimized systems exist, at much higher cost.

No such consumer/corporate demand exists yet for tele-immersion; however, the near-term ubiquity of 3D graphics engines, expected implosion of telecommunications costs, and emergence of new display technologies are reasons for timely experimental development of integrated systems. We hope to inspire private sector investment by describing prototypes of fully integrated tele-immersion hardware and software. Many of the barriers are market-based, but several are true technical research issues. Below, we identify and propose to address a set of these research issues.

The tele-immersion system of 2009 would ideally:

- Support one or more flat panels/projectors with ultra-high color resolution (say 5000x5000)
- Be stereo capable without special glasses
- Have several built-in micro-cameras and microphones
- Have tether-less, low-latency, high-accuracy tracking

- Network to teraflop computing via multi-gigabit optical switches with low latency
- Have exquisite directional sound capability
- Be available in a range of compatible hardware and software configurations
- Have gaze-directed or gesture-directed variable resolution and quality of rendering
- Incorporate AI-based predictive models to compensate for latency and anticipate user transitions
- Use a range of sophisticated haptic devices to couple to human movement and touch
- Accommodate disabled and fatigued users in the spirit of the Every Citizen Interface to the NII [2]

What we have as parts to integrate into 1999 systems are:

- Heavy, moderately expensive 3-tube projectors as the only straightforward stereo-capable projection devices
- Large projection distances needed for rear projection
- Medium resolution (1280x1024 pixel) displays
- Moderately awkward stereo glasses with limited view angle
- Stereo graphics hardware that integrates poorly with non-stereo camera input
- Imprecise electromagnetic tethered tracking with significant latency
- “Best effort” networking with random latency
- Expensive multi-processor workstations and rendering engines (\$25,000-\$200,000/screen)
- Primitive software models of user interactions within VR and tele-immersive systems
- Very primitive hardware devices for haptic interaction

In addition to the obvious dependency on improvements in display devices, computing hardware, and network integration, the tele-immersion system of 2009 will need to rely on emerging results from the computer science research community, including specifically:

- Data intensive computing and data mining
- Image-based modeling
- Digital audio/video transmission
- Recording/playback of sessions
- Every citizen interfaces
- Gesture, speech, and gaze interaction

2.C. Challenges of Tele-Immersion

Tele-immersion has emerged as a high-end driver for the Quality of Service (QoS), bandwidth, and reservation efforts envisioned by the NGI and Internet2 leadership. From a networking perspective, tele-immersion is a very challenging technology for several reasons:

- The networks must be in place and tuned to support high-bandwidth applications
- Low latency, needed for 2-way collaboration, is hard to specify and guarantee given current middleware
- The speed of light in fiber itself is a limiting factor over transcontinental and transoceanic distances
- Multicast, unicast, reliable and unreliable data transmissions (called “flows”) need to be provided for and managed by the networks and the operating systems of supercomputer-class workstations
- Real-time considerations for video and audio reconstruction (“streaming”) are critical to achieving the feel of telepresence, whether synchronous or recorded and played back
- The computers, too, are bandwidth limited with regard to handling very large data for collaboration
- Simulation and data mining are open-ended in computational and bandwidth needs—there will never be quite enough computing and bits/second to fully analyze, and simulate reality for scientific purposes

2.C.1. Tele-Immersion Flow Types

Progress in all these areas, however, is expected; tele-immersion serves as an integrating technology as pieces of the solution are contributed by the community and computer/networking industry. The following table, developed in discussions with Rick Stevens, director of the Math and Computer Science Division at Argonne National Lab, gives

our current best *estimations* and *opinions* of the attributes of the nine types flows simultaneously needed for an n-way compute and data-intensive audio, video, and haptic (touch) tele-immersive session. [35] The research agenda for the coming years *very much* involves validating this table and creating software intelligence to compensate for the otherwise unachievable.

Tele-Immersion Data Flow Types							
Type	Latency	Bandwidth	Reliable	Multicast	Security	Streaming	DynQoS
Control	< 30 ms	64Kb/s	Yes	No	High	No	Low
Text	< 100 ms	64Kb/s	Yes	No	Medium	No	Low
Audio	< 30 ms	Nx128Kb/s	No	Yes	Medium	Yes	Medium
Video	< 100 ms	Nx5Mb/s	No	Yes	Low	Yes	Medium
Tracking	< 10 ms	Nx128Kb/s	No	Yes	Low	Yes	Medium
Database	< 100 ms	> 1GB/s	Yes	Maybe	Medium	No	High
Simulation	< 30 ms	> 1GB/s	Mixed	Maybe	Medium	Maybe	High
Haptic	< 10 ms	> 1 Mb/s	Mixed	Maybe	Low	Maybe	High
Rendering	< 30 ms	>1GB/s	No	Maybe	Low	Maybe	Medium

The columns represent flow-type attributes:

- **Latency** is the sum of all delays in the system, from the speed of light in fiber, to operating system overhead, to tracker settling time and screen refresh
- **Bandwidth** is the bits/second the system can transmit
- **Reliable** flows are verified and retransmitted if bad
- **Multicast** flows go to more than one site at once
- **Security** involves encryption overhead that may or may not be warranted or legal
- **Streaming** data is a constant flow of information over time, as with video, audio and tracking
- **Dynamic QoS** can provide ways to service bursty high-bandwidth needs on request

The rows indicate the data flow types:

- **Control information** consists of data that is used to manage the tele-immersion session, to authenticate users or processes, to launch processes, to control the display or tracking systems, and to communicate out of band between the world servers and VR systems.
- **Text** provides simple communications capability within collaborative sessions for simple note taking and passing. Text can also command Unix processes driving the environments.
- **Audio** gives ambient auditory cues, allows voice communications among users, and is used to issue commands via voice recognition and speech synthesis. A typical application may use multiple audio streams.
- **Video** can allow teleconferencing or remote monitoring displayed within the virtual world. Synthetic 2D animated bitmaps in video format have application as well.
- **Tracking** is achieved with location and orientation sensors, and captures the position and orientation of the user. Typically this data is streamed to the computer responsible for computing the perspective of the scene. Tele-immersion requires tracking data to be shared among sites. Most VR systems only head and hand track; future systems will have many more sensors to track more complex posture and body motions.
- **Database** is the heart of a tele-immersion application world. The database contains the graphical models of virtual scenes, objects, and data, and since the database is used to provide the models that are rendered, it must be maintained in a coherent state across multiple sites. Databases might be as simple as shared VRML files or as complex as multi-terabyte scientific datasets, VR extensions of video serving, or even Virtual Director recorded sessions. (Virtual Director is a joint EVL/NCSA development project. [29])
- **Simulation** provides the basis for dynamic behaviors, like responding to the users' actions. Small-scale simulations often run on the computer also generating the VR experience, but frequently the simulation will need a dedicated supercomputer. [28] User input is captured and transmitted to the simulation via the network and the simulation will generate an update, which is then propagated to each user site for local

rendering. Typically the data transferred to the simulation is considerably smaller than the data returned by the simulation. For example, if the user is conducting an interactive molecular docking experiment, only tracking data needs be sent to the molecular model indicating the location of the user's hand, however, in response, the simulation will return updated coordinates of hundreds or thousands of atoms.

- **Haptics** include force and touch sensing/feedback devices and use a variety of sensors and actuators that are “attached” to the hands and/or legs of users. Some systems now generate haptic “images” that augment or replace visual images (e.g., a user feeling the magnetic field around a star simulation or perceiving the texture of an atomic scale surface being imaged by a scanning microscope). Haptics are particularly sensitive to latency and jitter (instantaneous variations in latency).
- **Rendering** is the transformation of geometric information (polygonal or volumetric) into images for display. All VR environments primarily render graphics locally. As networks provide bandwidth adequate for compressed HDTV, however, it will become reasonable and efficient for scenes to be rendered remotely and transmitted to each site in real time.

Note that large peak transfer rates for database, simulation, and rendering are due to the fact that relatively simple actions in the virtual world by a user can cause a considerable demand for synchronization or consistency updates at each participating site. Real-time rendering requirements may imply the need to distribute updates within one frame update interval (1/30-1/60 second) to avoid jerkiness or pauses in the graphics or inconsistencies in the shared world. While intelligent and speculative pre-fetching can often reduce the need for peak bandwidth, the ultimate limit is the nature and complexity of the world model and the expectations of the user.

2.C.2. Lag vs. Network Latency

In tele-immersion systems, an important distinction must be made between the notions of network performance and user perceived lag in the virtual environment. “Lag” is the term used to describe the perceived sum of all the sources of latency in a system.

Typically, it is thought of as the delay between action in the real world (e.g., as captured by tracking or haptics) and the perceived response of the system to that action. Lag is the critical issue for usability; reducing lag is a major technical challenge. Communications latency is only one component of tele-immersion lag. Effective solutions to reducing lag must attack the component sources of latency at all levels of the system. VR system lag is the result of delays in rendering, display, tracking, simulation, communications, and synchronization. There are multiple sources of latency in the communications system alone:

- Transmission latency—the time it takes to send a packet from one node to another
- Bandwidth or transfer latency—the time it takes to move data due to the size of the transfer
- Switching or routing latency—the sum of the delays due to the fact that the network is not composed of just point-to-point links
- Contention—delay caused by competition for limited resources in the network (bandwidth, queues, etc.)
- Protocol latency—due to the segmentation and re-assembly operations to build data packets and the header processing for protocol stacks.

Most users have difficulty manipulating objects in VR once lag exceeds 200ms. When the virtual display is coupled with the real world, as in tele-robotics, this limit is approximately 30ms. Non-network components of the VR system often together exceed 200-300ms, so there is actually very little room for wide-area communications delay in the lag budget. Research into asynchronous tele-immersion models should improve this situation; however, absolute limits of transmission latency due to time-of-light round trips may ultimately limit the geographical extent of tightly coupled tele-immersion environments.

2.C.3. Quality of Service (QoS)

QoS requirements need the ability to:

- Assign minimum service guarantees and relative priorities to each of many streams
- Specify notification and compensation actions if the QoS dynamics of the network change over time
- Predict the reliability of service guarantees or service estimates

Often the application can make some intelligent use of QoS information provided by the network or middleware software layer by taking corrective or compensatory actions during execution, provided dynamic control interfaces are available. Simplified QoS concepts, such as differentiated services, allow for prioritized traffic without specific service guarantees; this may be a way to provide uncongested bandwidth for tele-immersion and other collaborative users that is scalable.

3. Correspondences/Dependencies: Tele-Immersive Device Design Concepts

3.A. Motivation for Desktop/Office-Sized VR Display Devices

The VR community needs to conduct research in third-generation VR devices to construct software-compatible, variable-resolution and desktop/office-sized prototypes and products, which evolve over time as technology improves and needs become more demanding; the hardware now needs to be made smaller, higher resolution, and more adaptable to the human and his/her workspace. The recommendation is to procure, evaluate, and integrate a variety of emerging display devices, such as large color plasma displays, LCD projectors, LED panels, Digital Light Valves (DLVs), Grating Light Valves (GLVs), and Digital Micro Mirror Displays (DMDs).

To construct the tele-immersive office workspace, one would want affordable wall-sized high-resolution border-less displays with low lag and undiminished image intensity when viewed at an angle. Given that such a display does not exist today, we must rather learn from assembling new VR systems from available components.¹ We must push screen technology development by creating pressure from the computational science and engineering communities with compelling applications projects.

We describe several devices, each of which addresses different major issues in the tele-immersion/VR human computer interface:

- **ImmersaDesk3 Plasma Panel Desktop VR**
- **Personal Augmented Reality Immersive System (PARIS)**
- **Personal Penta Panel (P3)**
- **Totally Active Workspace (TAWs)**
- **CyberCeiling**
- **CAVEscope**

3.B. New Immersive Display Technologies

In the context of building new VR devices, the community needs to investigate the viability, flexibility of operation and breadth of application of the following new display technologies as compared to current 3-tube projector systems:

- *Liquid Crystal Display (LCD) projectors and panels.* They are achieving better resolution now (1280x1024), but have too high lag to be used for stereo unless two projectors are used with shutters. [www.nec.com, www.angleview.com, www.electrohome.com]
- *Digital Micro-mirror Displays (DMDs).* These are good resolution (1280x1024), and theoretically fast enough for stereo, but the supplied firmware does not support stereo. [www.electrohome.com, store.infocus.com/lp420/index.htm, www.ti.com/dlp/docs/business/manufacturers/, www.ti.com/dlp/docs/business/resources/press/pr94/448asc.htm, www.ti.com/dlp/docs/business/manufacturers/ask1.html]
- *Plasma panel displays.* These are low-medium resolution (800x480) but probably fast enough to do stereo with the proper driver electronics. These displays have electronics mounted around their edges that make border-less multi-screen configurations a challenge. [www.nec.com, www.fujitsu.com]
- *Light Emitting Diode (LED) displays.* These are low resolution right now (e.g., 208x272 and 320x192) but bright and border-less, in principle. [www.lumagraph.com/lumatile.html, www.daktronics.com/PAGES/Prostar.HTM, www.microweb.com/kwu/fdisplay.html]

¹ Several companies, like Panoram [www.panoramtech.com/realitycenter/wallandesk.html] and VRex [www.vrex.com], offer well-designed, non-tracked displays for the office and showroom. Barco and Fakespace have products similar to the ImmersaDesk. The goal of this paper is not to suggest competition with the commercial sector, but to investigate and inspire new display and control technologies for the human-centered interface to tele-immersion.

- *Digital Light Valve (DLV) displays.* These new desktop projection displays have latency problems for stereo use; they can switch fast enough but do not go to black in the required time. A 2Kx2K resolution version has been built [www.almaden.ibm.com/journal/rd/423/melcher.txt]
- *Grating Light Valve (GLV) displays.* Recently demonstrated in prototype form, this laser-driven micro-electromechanical display is capable of HDTV resolution at 96Hz, very promising for VR. Switching speeds are extremely low, allowing a linear array of deflectable ribbon picture elements to scan out an image [26] [www.siliconlight.com]

3.C. ImmersaDesk3

3.C.1. Plasma Panel Desktop Device—A Design Exercise

ImmersaDesks and Responsive Workbenches are large because the available 3-tube projection technology has a limit to how small the screen can get (approximately 6' diagonal). Rear projection distances are significant, even when folded with mirrors, and the projector itself is quite large and heavy. Both of these devices are sized for a laboratory, and are too large for a typical faculty office or cubicle. We built a prototype device, called the ImmersaDesk3 to test the plasma panel technology currently available at 640x480 resolution for US\$10,000.

The ImmersaDesk3 is configured so a user can position the screen at any angle from horizontal to vertical, forward or back, on the desk. The angle can be measured automatically so that the correct perspective view of the computer-generated images for the tracked user is presented. Cameras can be added to this configuration to make image/gesture recognition, tether-less tracking and tele-immersion experiments possible. Given its configuration flexibility, the ImmersaDesk3 is also amenable to the integration of haptic (tactile input/output) devices.

We built our system around the Fujitsu PDS4201U-H Plasmavision display panel. The Plasmavision has an active display area of 36x20 inches (in a 16:9 aspect ratio); the entire panel is 41x25x6 inches and weighs 80 pounds.

We mounted the Plasmavision on a modified office desk.

To accommodate different applications and for greater flexibility, we wanted to be able to position the screen vertically (perpendicular to the desktop), horizontally (flat on the desktop), or at an angle in between. The panel is too heavy for users to shift easily, so we mounted it on hydraulic supports with a hand crank to adjust the angle.

3.C.2. Problems Encountered with Plasmavision Plasma Panel Displays

The Plasmavision outputs a 30Hz, interlaced, NTSC resolution image. When this is used for a stereoscopic display, each eye is only seeing a 30Hz signal, and the flicker is very noticeable; prolonged exposure can give many users headaches. Using the NTSC field-interleaved format for stereo yields only 640x240 pixel resolution for each eye's image. We also found that the red and green phosphors do not decay quickly enough. When we look at a stereo test pattern, which displays separate red, green, and blue color bars for each eye, only the blue bar is sufficiently extinguished; the red and green bars are still visible to the *wrong* eye. In an informal test of 16 users at SIGGRAPH '98, we noted that 25% of them could fuse the full-color images, while 50% could only fuse the images when the red and green channels were disabled, so that the images were just shades of blue; 25% of the users could not fuse the images at all. [25]

The Plasmavision is electromagnetically noisy, so much so that it interferes with the accuracy of magnetic tracking systems.

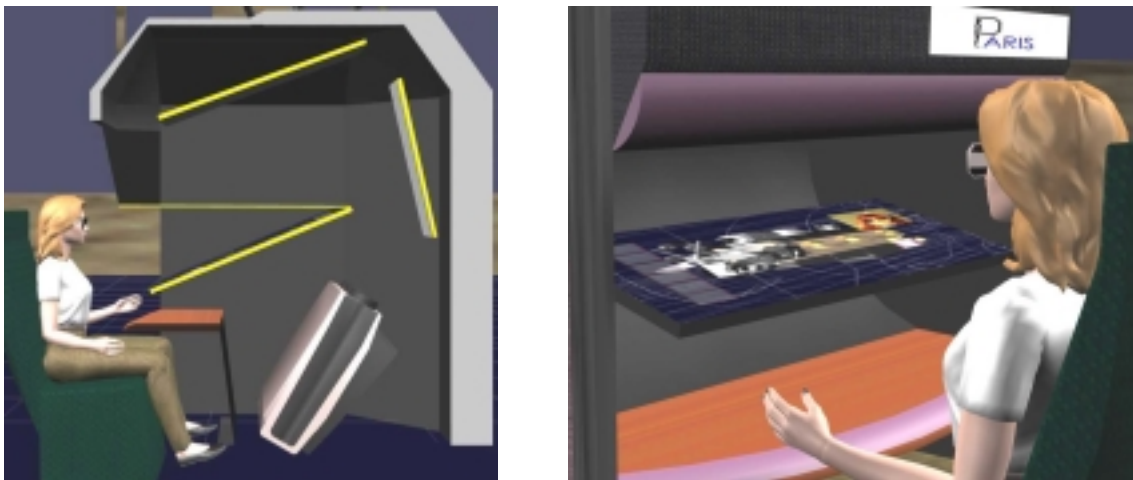
of view than a conventional monitor, yet it fits well on a desktop, something a 42" monitor or projection system cannot do.

Current plasma panel technology has severe limitations as a stereo display device for projection-based VR systems. The inability to easily sync the plasma panel to shutter glasses and the red/green phosphor decay problem preclude clear stereo. The low resolution and 30Hz frame rate also prevent current panels from being serious contenders in this field. Although flat panels can significantly save space, larger display systems would need larger panels or tiled panels. Current plasma panels have borders that prevent seamless tiling.

Nevertheless, the concept of a wide-field-of-view, desktop VR system and space-saving flat panel technology for CAVEs and other large displays is appealing. We look forward to improvements in flat panel technology as it evolves.

3.D. Personal Augmented Reality Immersive System (PARIS)

Twenty years ago, Ken Knowlton created a see-through display for Bell Labs using a half-silvered mirror mounted at an angle in front of a telephone operator. The monitor driving the display was positioned above the desk facing down so that its image of a virtual keyboard could be superimposed on an operator's hands working under the mirror. The keycaps on the operator's physical keyboard could be dynamically relabeled to match the task of completing a call as it progressed. Devices that align computer imagery with the user's viewable environment, like Knowlton's, are examples of *augmented reality*, or *see-through VR*. More recently, researchers at the National University of Singapore's Institute of Systems Science built a stereo device of similar plan using a Silicon Graphics' monitor, a well-executed configuration for working with small parts in high-resolution VR [15]. Neither of these systems provides tracking, but rather assumes the user to be in a fixed and seated position.



Artist renderings of PARIS system: Cut-away side view of PARIS system (left) and over-the-shoulder view (right). Image courtesy Sam Thongrong, Electronic Visualization Laboratory, University of Illinois at Chicago, 1999.

We are currently building a desktop VR device, the Personal Augmented Reality Immersive System (PARIS), a third-generation version of this concept. We insure that a keyboard is integrated, and that suitably positioned cameras can capture facial expressions and head position. Gesture recognition can come from tracking, as well as processed camera input. Audio support can be used for voice recognition and generation as well as for recording and tele-immersion sessions. (See "Multimodal Human Computer Interfaces" sidebar below.)

Since we are committed to stereo in PARIS, and would like as high as possible resolution, we cannot initially use a plasma panel display for the reasons belabored in the previous section. Instead, we will use two 1280x1024 LCD projectors with electronic shutters compatible with active glasses to achieve stereo separation²

² LCDs have very high lag so time-based stereo separation is not possible with a single projector; instead, we propose two projectors with external blinking shutters to separate the left and right eye views. VRex, Inc. markets a line of LCD stereo

We can also use PARIS to prototype passive (polarized) stereo since we can polarize the two projector outputs, allowing very inexpensive and lightweight glasses to be incorporated, an important feature for use in museums and schools. If plasma or LED panel displays ever have excellent brightness, stereo speeds, and high-resolution, these would be preferable devices to adapt.

Multimodal Human Computer Interfaces

Thomas S. Huang

Introduction

As information systems become global, there is increasing interest in making access to computing, communications, and information repositories universal. To achieve this goal, it is imperative to explore the use of new modalities (such as vision and speech) in human-computer interfaces (HCI). The Image Laboratory at the Beckman Institute, University of Illinois at Urbana-Champaign, has been collaborating with EVL to develop multimodal human computer interfaces for existing (ImmersaDesk, CAVE) and future (PARIS, etc.) virtual environments. [30, 31, 32, 33, 34]

Current Research

Current research projects in Multimodal HCI include:

1. Integrating speech and vision-based hand gesture recognition in display controls in virtual environments. Algorithms have been implemented on ImmersaDesks to use the hand to point to a specific object in the display while using voice to ask questions about the object, or to use the hand and voice together to navigate through 3D terrain (the hand is used to steer and the voice is used to stop, go forward, change speed, etc.).
2. Real-time algorithms for detecting, recognizing, and tracking faces and recognizing (together with tone of voice) emotion.
3. Use of visual lip-reading to help the accuracy of audio speech recognition in noisy environments.
4. We have started to work on the difficult problem of tracking fingers visually, which will be needed for virtual manipulation in the PARIS environment.

Challenging Research Issues

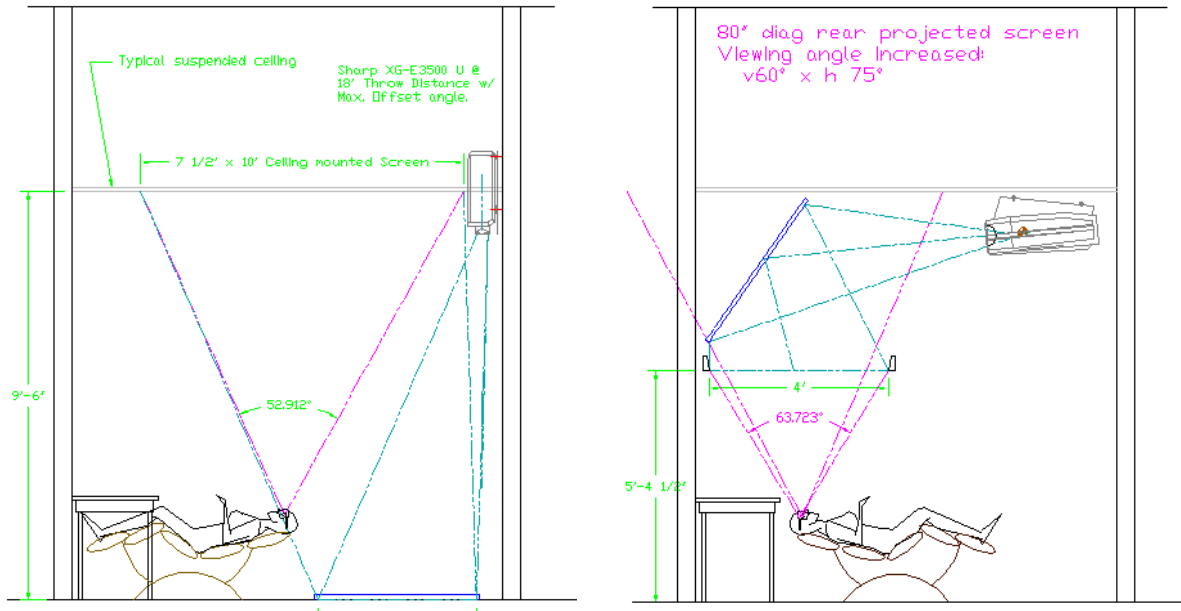
In addition to the visual tracking of hand/fingers, some of the major research issues for the future are:

1. Visual tracking of humans in less constrained situations. Current work on visual human tracking concentrates on tracking isolated human parts (face, arm/hand, body) and imposes severe constraints, such as front view for faces, long-sleeve shirts for arm/hand, and tight clothing for body. In many future applications, it will be necessary to track the entire human robustly, in real time, and without constraints.
2. Fusion of cues from different modalities. What are good architectures for integrating the information from different modalities to reach decisions?

3.E. CyberCeilings, designed for the Last Unused Projection Surface

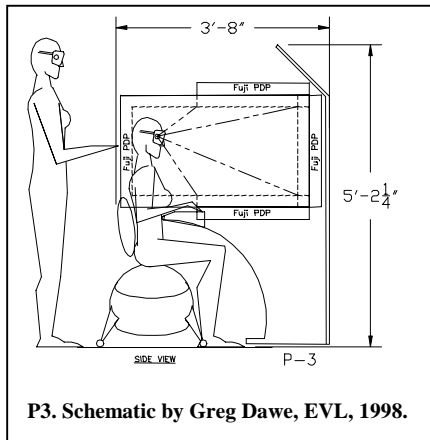
In trying to fit large screens and VR into offices, use of overhead space or the ceiling is conceivable, and has distinct advantages in hospital patient settings, assuming the room can be made dark enough. The drawings below indicate some options for ceiling-mounted front projection with a mirror on the floor, and a smaller, rear projection overhead display. Different lensing can alter the projection distances in the former example. The chair shown is a commercially available executive motorized recliner, but could be replaced by a bed in a hospital setting. This configuration has the benefit that the user may not need to be tracked since body position is fixed and head rotation is not accounted for in projected VR environments.

projectors that use polarization multiplexing with a faceplate over the LCD to separate the left and right eye images. The effective resolution is halved since the displays are spatially multiplexed rather than time multiplexed, but there is evidence that the brain re-integrates the information, lessening the problem.



Schematic of CyberCeiling. Image courtesy of Greg Dawe, Electronic Visualization Laboratory, University of Illinois at Chicago, 1999.

3.F. Personal Penta Panel (P3) or Dilbert's Dream



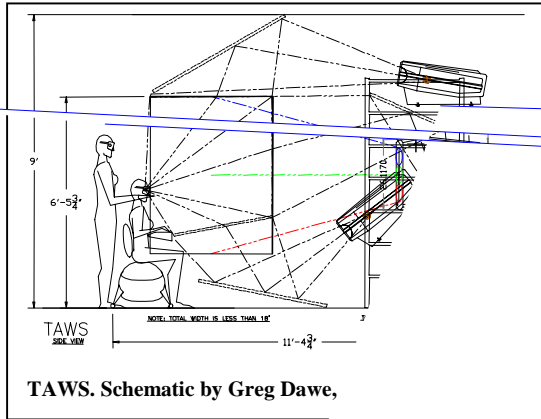
P3. Schematic by Greg Dawe, EVL, 1998.

Dilbert's Dream is conceived as an office cubicle whose walls and desk are made from border-less stereo-capable high resolution panels, not, unfortunately, obtainable in the current millenium.

Alternatively, we are proposing a "desktop" cubicle. The Personal Penta Panel (P3) is a box made out of 42" diagonal plasma panels. The user places his/her tracked head and hands into the box of screens and is presented with a surround (non-stereo) view. Each panel would have a frame around it, creating seams between screens that would be difficult to eliminate. There are, however, optical methods to relay an image a few inches forward, which could be used to (mostly) eliminate the effects of the frames. Such a device would be useful for all but very close viewing, even in non-stereo, as we wait for the needed technological improvements in panels.

Scott Adams, creator of Dilbert, recently suggested that this kind of device may be harmful to programmers! In his article "Gene Fool" in Time Magazine, [27] he explains: *"But unlike the sterile boxes of today, every cubicle will be a technology wonderland customized for the occupant. Flat-panel screens on each wall will give the impression you are in a hot air balloon floating over the Alps. Noise-cancellation technology will block out the surrounding sounds while providing a symphony within the cubisphere. The computer will continue its evolution to a full entertainment center, providing a constant supply of first-run movies, live nudity, gambling and video conferencing. The engineer's chair will be soft and warm, conforming to the body and providing simulated motion and vibration to match the entertainment. The cubicle experience will be so much better than life on the outside, engineers won't want to leave. That could be a problem. I heard about an experiment where rats were given the choice between food and cocaine. They chose the cocaine until they starved. The same thing will happen to the engineers. I predict they'll all starve to death inside their cubicle wonderlands. I just hope no one blames me."*

3.G. Totally Active Work Space (TAWs)

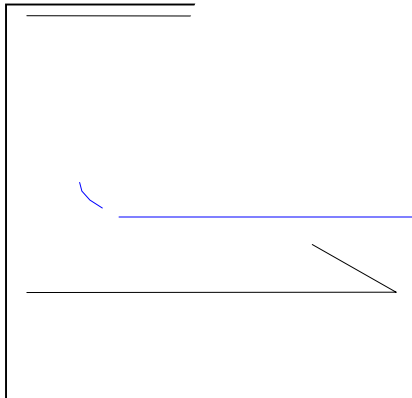


Pending the availability of suitable plasma, LCD, or LED panels, we have built screens into a rear-projection desktop structure to simulate the Totally Active Work Space (TAWS)—the ultimate Dilbert’s Dream or *cubisphere*. TAWS is large

TAWS. Schematic by Greg Dawe,

3.H. CAVEscope: Sim

All projection-based VR d narrow angle, the ~five d resolution displays that, degree angle of view. Ir feet, a *rather less-than* but resolution itself c scanning fluctuation however, provides



desired. The CAVEscope would be tracked so that it would present the proper perspective projection. Touch screen technology could also be available for user input. A miniature television camera mounted on the CAVEscope could enable tele-conferencing. Users can see and talk to each other using CAVEscopes, or position their devices for coverage relevant to the task at hand. CAVEscope combines the intuitive navigational

capabilities of projection-based VR with the detailed view of the LCD portal, all under user control.

CAVEscope should also be usable in an office setting with front projection VR on the office walls and desktop, such as has been proposed by the Advanced Network and Services-sponsored National Tele-Immersion Initiative. [\[www.advanced.org/teleimmersion.html\]](http://www.advanced.org/teleimmersion.html) Since the wall projections are used mainly for navigation and context, not for detail work, the quality of the projected images could be less than optimal, as long as the CAVEscope image is suitably bright and sharp.

Since LCD panel technology does not permit Crystal Eyes-type stereo (due to high lag) at this point, we will need to work with a mono image, pending the availability of a compatible stereo-capable panel in the future. Taking the stereo glasses on and off is an annoyance.

³ Some flight simulators have elaborate mechanisms to create high-resolution images at the pilot’s center of interest by using a second projector inset at higher resolution. Since VR users have much more freedom than a pilot to move and look around, this technique will not work well since the inset projector, whose image is usually positioned by a moving mirror, has a limited range of motion and focus.

Tracked hand-held panels have been suggested as portals into virtual and augmented reality spaces for some time, although, on videotape, the concept is simulated with chroma keying. [17] Discovering where to look in virtual space is a large part of the problem with narrow-angle-of-view devices like panels held at arms length, VR binoculars, or even head-mounted displays. CAVEscope affords the user both the navigational and wide field of view of projection-based VR with a real-time high-resolution inspection capability. Since CAVEscope has its own rendering engine, the software can be tuned to provide much more detailed rendering in the designated area of interest, which could even be behind or above the user where there are no projected screens!

In addition, the user can easily enough freeze the motion and build up the display or re-render it with ray tracing, a type of successive refinement not normally usable in VR. We believe these notions will provide enhanced performance in accuracy, resolving power, flexibility of operation, user friendliness and navigation for scientists and engineers using projection-based VR for discovery and observation.

4. Questions for the Future

1. It is appropriate to stipulate that market forces for commodity computing improvement will not be driven by virtual reality or tele-immersion needs for at least a few years. However, it is clear that designers of future workspaces will use large, high-resolution flat displays given availability at almost any cost. How do we motivate the creation and marketing of these high-resolution panel devices most effectively? How do we convince the panel designers and manufacturers to incorporate features that support stereo?
2. It is also clear that 3D real-time graphics generation suitable for VR is achievable cheaply from the example set by the Nintendo64 Silicon Graphics chip set. How do we get such displays boosted in resolution and enable, say, dual-channel output for cheap dual projector displays, or stereo capability for the appropriate panel displays? How do we encourage manufacturers of micro electromechanical displays like the GLV and DLV to develop stereo-capable firmware?
3. Audio is well handled as (one-dimensional) phone calls; video and audio are supported by teleconference (two-dimensional) calls. How do we push the development of operating systems and networks that handle tele-immersion (three-dimensional) calls? Do we develop personal switches and /or routers that have graphics/audio/video boards, or will the personal computer industry develop machines that handle multiple types of network flows?
4. How do we motivate the development of haptic devices that do not injure users from repetitive use? How do we integrate video and motion detection to have our computers recognize our needs and wants? How do we send haptic inputs and haptic displays over networks?
5. How do we build systems sufficiently rich in I/O to accommodate every citizen, including the elderly? Will we succeed in standardizing/propagating these interfaces before we get too old to enjoy them?
6. What are the implications of transcontinental and transoceanic distances for tele-immersion.
7. How do we archive tele-immersive sessions in ways we can play back and/or edit in the near and far future?

5. Acknowledgments

EVL's virtual reality research, collaborations, and outreach programs are made possible by major funding from the National Science Foundation (NSF), awards EIA-9802090, EIA-9871058, EIA-9720351, ANI-9712283, and ACI-9418068, as well as NSF Partnerships for Advanced Computational Infrastructure (PACI) cooperative agreement ACI-9619019 to the National Computational Science Alliance. EVL also receives support from the US Department of Energy Accelerated Strategic Computing Initiative (ASCI) and Pacific Interface on behalf of NTT Optical Network Systems Laboratory in Japan.

The CAVE and ImmersaDesk are trademarks of the Board of Trustees of the University of Illinois.

6. References Cited

- [1] D. Cox, "Cosmic Voyage: Scientific Visualization for IMAX film" (Animator's Sketches), and "Cosmic Voyage: Galaxy Formation and Interaction" (Computer Animation Festival), SIGGRAPH 96 Visual Proceedings, Computer Graphics Annual Conference Series ACM SIGGRAPH, 1996, pp. 129 and 174.
- [2] More Than Screen Deep: Toward Every-Citizen Interfaces to the Nation's Information Infrastructure, Computer Science and Telecommunications Board, National Research Council, National Academy Press, 1997.
- [3] M. Czernuszenko, D. Pape, D. Sandin, T. DeFanti, G.L. Dawe, and M.D. Brown, "The ImmersaDesk and Infinity Wall Projection-Based Virtual Reality Displays," Computer Graphics, ACM SIGGRAPH, Volume 31, Number 2, May 1997, pp. 46-49. [<http://www.evl.uic.edu/pape/CAVE/idesk/paper/>]
- [4] T.A. DeFanti, M.D. Brown and R. Stevens (Guest Editors), "Virtual Reality Over High-Speed Networks," IEEE Computer Graphics & Applications, July 1996, Vol. 16, No. 4, pp. 14-17, 42-84.
- [5] T.A. DeFanti, I. Foster, M. Papka, R. Stevens, and T. Kuhfuss, "Overview of the I-WAY: Wide Area Visual Supercomputing," International Journal of Supercomputer Applications and High Performance Computing, Vol. 10, No. 2/3, Summer/Fall 1996, pp. 123-131.
- [6] M. Ghazisaedy, D. Adamczyk, D.J. Sandin, R.V. Kenyon, and T.A. DeFanti, "Ultrasonic Calibration of a Magnetic Tracker in a Virtual Reality Space," Proceedings of the IEEE Virtual Reality Annual International Symposium (VRAIS 95), Research Triangle Park, NC, March 1995.
- [7] A. Johnson, J. Leigh, J. Costigan, "Multiway Tele-Immersion at Supercomputing '97," IEEE Computer Graphics and Applications, July 1998.
- [8] R. Kenyon and M. Afenya, "Training in Real and Virtual Environments," Annals of Biomedical Engineering: Starkfest Conference Proceedings, Biomedical Engineering Society, Vol. 23, 1995, pp. 445-455.
- [9] C. Kesselman, R. Stevens, T. DeFanti, J. Bannister, I. Foster, "Tele-Immersive Applications on the Next-Generation Internet," (white paper), Workshop on Research Directions for the Next Generation Internet, Computing Research Association, sponsored by the Federal Large Scale Networking Working Group (LSN) of the National Science and Technology Council's Committee on Computing, Information, and Communications R&D Subcommittee, May 13-14, 1997.
- [10] S. Latta, "Creating Computer Graphics Choreography with Virtual Director," NCSA Access, Summer 1996, Volume 10, Number 2, National Center for Supercomputing Applications, University of Illinois at Chicago, p. 5, [<http://www.ncsa.uiuc.edu/Pubs/access/96.2/VirtualDir.html>]
- [11] V.D. Lehner and T.A. DeFanti, "Distributed Virtual Reality: Supporting Remote Collaboration in Vehicle Design," IEEE Computer Graphics & Applications, March/April 1997, pp. 13-17.
- [12] J. Leigh, T.A. DeFanti, A.E. Johnson, M.D. Brown, and D.J. Sandin, "Global Tele-Immersion: Better Than Being There," ICAT '97, 7th Annual International Conference on Artificial Reality and Tele-Existence, December 3-5, 1997, University of Tokyo, Japan, Virtual Reality Society of Japan, pp. 10-17.
- [13] J. Leigh, A. Johnson, and T. DeFanti, "CAVERN: Distributed Architecture for Supporting Scalable Persistence and Interoperability in Collaborative Virtual Environments," Virtual Reality: Research, Development and Applications, Vol. 2, No. 2, December 1997, pp. 217-237.
- [14] J. Leigh, A.E. Johnson, T.A. DeFanti, "Issues in the Design of a Flexible Distributed Architecture for Supporting Persistence and Interoperability in Collaborative Virtual Environments," SC'97 Proceedings, Sponsored by ACM SIGARCH and IEEE Computer Society, November 15-21, 1997, CD ROM.
- [15] T. Poston, "The Virtual Workbench : Dextrous Virtual Reality," G. Singh, S.K. Feiner, D. Thalmann (eds.), Technical Report TR94-132-0, Institute of Systems Science, National University of Singapore, August 1994, [<http://www.iss.nus.sg/iss/techreport>]
- [16] H. Korab and M.D. Brown, (eds.), "Virtual Environments and Distributed Computing at SC'95: GII Testbed and HPC Challenge Applications on the I-WAY," a publication of ACM/IEEE Supercomputing '95, [www.ncsa.uiuc.edu/General/Training/SC95/GII.HPCC.html]
- [17] "Starfire: A Vision of Future Computing," (videotape), Sun Microsystems, 1995. (Also in Tog on Software Design, Addison-Wesley, Fall 1995.)

- [18] J.E. Smith and F.W. Weingarten (eds.), *Research Challenges for the Next Generation Internet*, Computing Research Association, 1997, p. 20.
- [19] R. Stevens, P. Woodward, T. DeFanti, and C. Catlett, "From the I-WAY to the National Technology Grid," *Communications of the ACM*, Vol. 40, No. 11, November 1997, pp. 51-60.
- [20] M. Thiébaux, *Virtual Director*, MS Thesis, Electronic Visualization Laboratory, Department of Electrical Engineering and Computer Science, University of Illinois at Chicago, June, 1997.
- [21] L. Petrovich, K. Tanaka, D. Morse, N. Ingle, J. Morie, C. Stapleton, M. Brown, (eds.) "Visual Proceedings," *Computer Graphics, Annual Conference Series 1994, ACM SIGGRAPH, 1994*, pp. 219-264, [www.evl.uic.edu/EVL/VROOM/HTML/OTHER/HomePage.html]
- [22] "Virtual Reality at the University of Illinois 5/97," (videotape), Electronic Visualization Laboratory, University of Illinois at Chicago.
- [23] M. Brown, T. DeFanti, M. McRobbie, A. Verlo, D. Plepys, D.F. McMullen, K. Adams, J. Leigh, A. Johnson, I. Foster, C. Kesselman, A. Schmidt, S. Goldstein, "The International Grid (iGrid): Empowering Global Research Community Networking Using High Performance International Internet Services," *INET '99*, San Jose, June 22-25, 1998, San Jose, California, (accepted for publication).
- [24] McMullen, Donald F., Michael A. McRobbie, Karen Adams, Douglas Pearson, Thomas A. DeFanti, Maxine D. Brown, Dana Plepys, Alan Verlo, Steven N. Goldstein. February 18-20, 1999 (accepted for publication). *The iGrid Project: Enabling International Research and Education Collaborations through High Performance Networking. Internet Workshop '99 (IWS'99)*, Osaka, Japan.
- [25] Pape, Dave, Josephine Anstey, Mike Bogucki, Greg Dawe, Tom DeFanti, Andy Johnson, Dan Sandin, "The ImmersaDesk3—Experiences With A Flat Panel Display for Virtual Reality," *Proceedings of the Third International Immersive Projection Technology (IPT) Workshop*, Stuttgart, Germany, May 10-11, 1999, to be published by Springer Verlag, (accepted for publication).
- [26] D.M. Bloom, *The Grating Light Valve: Revolutionizing Display Technology*, SPIE Vol. 3013, Feb. 1997, pp165-171.
- [27] Scott Adams, "Gene Fool," *TIME Magazine*, December 7, 1998, pp. 216-217.
- [28] T.M. Roy, C. Cruz-Neira, T.A. DeFanti, D.J. Sandin, "Steering a High Performance Computing Application from a Virtual Environment," *Presence: Teleoperators and Virtual Environments*, Vol. 4, No. 2, MIT Press, Spring 1995, pp. 121-129.
- [29] Dan Reed, Roscoe Giles, and Charles E. Catlett, "Distributed Data and Immersive Collaboration," *Communications of the ACM*, November 1997, pp. 39-48.
- [30] V.I. Pavlovic, R. Sharma, and T.S. Huang, "Visual Interpretation of Hand Gestures for Human Computer Interaction: A review," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol. 19, No. 7, July 1997, pp. 677-695.
- [31] A.J. Colmenarez, B.J. Frey, and T.S. Huang, "A Probabilistic Framework for Embedded Face and Facial Expression Recognition," *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, June 23-25, 1999, Ft. Collins, CO.
- [32] M.T. Chan, Y. Zhang, and T.S. Huang, "Real-Time Lip Tracking and Bimodal Continuous Speech Recognition," *Proc. IEEE Workshop on Multimedia Signal Processing*, Dec. 7-9, 1998, Los Angeles, CA.
- [33] L.S. Chen, H. Tao, T.S. Huang, T. Miyasato, and R.N. Nakatsu, "Emotion Recognition from Audiovisual Information," *Proc. IEEE Workshop on Multimedia Signal Processing*, Dec. 7-9, 1998, Los Angeles, CA.
- [34] V.I. Pavlovic, B.J. Frey, and T.S. Huang, "Time-series classification using Mixed-State Dynamic Bayesian Networks," *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, June 23-25, 1999, Ft. Collins, CO.
- [35] R. Stevens and T.A. DeFanti, "Tele-Immersion and Collaborative Virtual Environments," *The Grid: Blueprint for a New Computing Infrastructure*, I. Foster and C. Kesselman (eds.), Morgan Kaufmann Publishers, 1999, pp. 131-158.

7. Biography

Thomas A. DeFanti

Electronic Visualization Laboratory
University of Illinois at Chicago
851 S. Morgan St., Room 1120
Chicago, IL 60607-7053

(312) 996-3002
(312) 413-7585 fax
tom@uic.edu
www.evl.uic.edu

Thomas A. DeFanti, PhD, is director of the Electronic Visualization Laboratory (EVL), a professor in the department of Electrical Engineering and Computer Science, and director of the Software Technologies Research Center at the University of Illinois at Chicago (UIC). He is also co-associate director of the Grid Division at the National Center for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign, and a co-PI of the National Computational Science Alliance, one of two recipients of the NSF Partnerships for Advanced Computational Infrastructure (PACI) program. DeFanti has strong working collaborations with other UIC departments, as well as NCSA and the Mathematics and Computer Science division of Argonne National Laboratory. DeFanti is chair of the UCAID/Internet2 Applications Strategy Council and member of the UCAID Board of Trustees.

Current research interests include: high-performance networking, virtual environments, tele-immersion, digital libraries, scientific visualization, new methodologies for informal science and engineering education, paradigms for information display, distributed computing, optimization for scalable computing, sonification, human/computer interfaces, and abstract math visualization.

DeFanti is an internationally recognized expert in computer graphics. Credits include: use of EVL hardware and software for the computer animation produced for the first “Star Wars” movie; early involvement in video game technology long before videogames became popular; contributor and co-editor of the 1987 National Science Foundation-sponsored report “Visualization in Scientific Computing;” recipient of the 1988 ACM Outstanding Contribution Award; an appointment in 1989 to the Illinois Governor’s Science Advisory Board; University Scholar for 1989-1992; appointed an ACM Fellow in 1994; appointed one of several US technical advisors to the G7 GIBN activity in 1995; appointed in 1999 to Mayor Richard M. Daley’s Council of Technology Advisors; principal investigator of the STAR TAP initiative to provide a persistent infrastructure to facilitate the long-term interconnection and interoperability of advanced international networking; and, recognition along with EVL co-director Daniel J. Sandin for conceiving the CAVE virtual reality theater in 1991.

He has also been active in the ACM SIGGRAPH organization and in the ACM/IEEE Supercomputing (SC) conferences. Current and past activities include: secretary of SIGGRAPH (1977-1981); co-chair of the SIGGRAPH 79 conference; chair of SIGGRAPH (1981-1985); past chair of SIGGRAPH (1985-1989); editor of the “SIGGRAPH Video Review” video publication, which he founded in 1979; and, member of the SIGGRAPH 92 and SIGGRAPH 94 conference committees. He was information architect of the SC’95 conference, responsible for the I-WAY and GII Testbed activities, and was a member of the SC’97 Program Committee.