

The Future of the CAVE

Research Article

Thomas A. DeFanti^{1*}, Daniel Acevedo², Richard A. Ainsworth³, Maxine D. Brown⁴, Steven Cutchin², Gregory Dawe¹, Kai-Uwe Doerr¹, Andrew Johnson⁴, Chris Knox², Robert Kooima⁵, Falko Kuester¹, Jason Leigh⁴, Lance Long⁴, Peter Otto¹, Vid Petrovic¹, Kevin Ponto¹, Andrew Prudhomme¹, Ramesh Rao¹, Luc Renambot⁴, Daniel J. Sandin⁴, Jurgen P. Schulze¹, Larry Smarr¹, Madhu Srinivasan², Philip Weber¹, Gregory Wickham²

1 *California Institute for Telecommunications and Information Technology (Calit2), University of California San Diego (UCSD), San Diego, California, USA*

2 *King Abdullah University of Science and Technology (KAUST) Visualization Laboratory, Thuwal, Saudi Arabia*

3 *Ainsworth & Partners, Inc., Ridgeway, Wisconsin, USA*

4 *Electronic Visualization Laboratory (EVL), University of Illinois at Chicago (UIC), Chicago, Illinois, USA*

5 *Center for Computation & Technology (CCT), Louisiana State University, Baton Rouge, Louisiana, USA*

Received 2 July 2010; accepted 31 August 2010

Abstract: The CAVE, a walk-in virtual reality environment typically consisting of 4-6 3 m-by-3 m sides of a room made of rear-projected screens, was first conceived and built in 1991. In the nearly two decades since its conception, the supporting technology has improved so that current CAVEs are much brighter, at much higher resolution, and have dramatically improved graphics performance. However, rear-projection-based CAVEs typically must be housed in a 10 m-by-10 m-by-10 m room (allowing space behind the screen walls for the projectors), which limits their deployment to large spaces. The CAVE of the future will be made of tessellated panel displays, eliminating the projection distance, but the implementation of such displays is challenging. Early multi-tile, panel-based, virtual-reality displays have been designed, prototyped, and built for the King Abdullah University of Science and Technology (KAUST) in Saudi Arabia by researchers at the University of California, San Diego, and the University of Illinois at Chicago. New means of image generation and control are considered key contributions to the future viability of the CAVE as a virtual-reality device.

Keywords: CAVE • Computer-supported collaborative work (CSCW) • Graphics packages • Image displays • Immersive environments • Interactive environments • Sonification • Tele-immersion • Virtual reality • Scalable multi-tile displays

© Versita Sp. z o.o.

*E-mail: tdefanti@ucsd.edu

1. The Classic CAVE

1.1. Brief History of the CAVE

The “classic” CAVE™ (Figure 1) [1] is a cube-shaped virtual-reality (VR) room, typically 3 m-by-3 m-by-3 m in size, whose walls, floor and sometimes ceiling are entirely made of computer-projected screens. All participants wear active stereo glasses to see and interact with complex 3D objects. One participant wears a six degree-of-freedom location and orientation sensor called a tracker so that when he/she moves within the CAVE, correct viewer-centered perspective and surround stereo projections are produced quickly enough to give a strong sense of 3D visual immersion.

Projection-based VR systems, such as CAVEs, feature surround viewing (ideally fully-surround, but usefully at least 90° in two dimensions such that users do not see the edges of the display). They offer stereo visuals. And, they track the user to provide the correct scene perspective rendering in a continuous manner. Viewer-centered perspective and surround viewing distinguish VR systems from 3D movies.

The classic CAVE was conceived and designed in 1991 by Tom DeFanti and Dan Sandin, who at the time were professors and co-directors of the Electronic Visualization Laboratory (EVL) at the University of Illinois at Chicago (UIC). Many students and colleagues over the years contributed to CAVE software and hardware development, as seen in this paper’s many references. Human factors studies such as discussed in [2] have helped quantify the accuracy of perception by users in CAVEs.

The goal of room-sized VR is to help scientists and engineers achieve scientific insight, and also to create a medium worthy of use by teachers, fine artists, architects, and archaeologists. The most advanced classic CAVE to date, called *Cornea*, was installed in mid-2009 at the King Abdullah University of Science and Technology (KAUST) by Mechdyne Corporation and co-located with a new-generation of unique VR systems, designed by a team led by DeFanti, who is now at the California Institute for Telecommunications and Information Technology (Calit2) at University of California, San Diego (UCSD). Other manufacturers of CAVE-like systems include Barco, N.V., Eon Reality, Inc., and Visbox, Inc. [3] lists 57 CAVEs at universities.

CAVE participants see projected computer-generated stereo scenes but can also see their arms and bodies and can easily interact with one another. The classic CAVE uses *active stereo*; the rendering computers generate perspective views that are projected onto the walls for the left and right eyes of the primary participant in

synchrony with electronic shutter glasses. The active eyewear is driven transparent in front of the left eye when the left eye image is projected, and opaque otherwise; similarly, the right eye receives the right image. Images need to be generated at 96 Hz (or higher) so each eye gets a flicker-free display. This is the way the new (2010) consumer 3D HDTVs work if they use active shutter glasses. For successful active stereo, all screens must be synchronized so that each eye sees only the left or right stereo image on every screen at the same time, a requirement that is non-trivial as the number of screens and projectors increases; therefore, the underlying hardware and drivers must support synchronization.

The classic CAVE uses rear-screen projection for the walls so the viewer does not cast shadows on the screens; however, as is often the case, when the floor is projected down from the top, the projection creates a small shadow around participants’ feet. A CAVE with three walls and a floor minimally requires a 13 m-by-10 m space with a 4.5 m high ceiling. Six-sided CAVEs (like KAUST’s *Cornea* and Iowa State’s C-6, as well as others built in Gifu, Stockholm, and Stuttgart, for example) feature rear projection from every direction, requiring doubly higher ceilings and elaborate support structures, and bottom-projected floor screens made out of thick museum-style aquarium-grade acrylic plastic that can withstand the weight of 6–10 people. Projection-based CAVEs require significant (and otherwise mostly wasted) rear-projection space, projectors costing \$5,000–\$500,000/screen, projector maintenance/alignment, lamp replacement, significant power and cooling, specialized screen material and controlled lighting conditions, all of which severely limit their acceptance and adoption in everyday workspaces, public venues, and homes. The screen material itself must allow severe off-axis viewing without objectionable intensity attenuation and, should minimize internal light spillage and reflection. (Head-mounted displays, which are ~1–2 inch screens mounted in some way in front of one’s eyes, are the earliest [4] technology for single-user VR, and are well suited for in-the-field and augmented (see-through) use. They will not be further discussed here since this paper focuses on future room-scale VR systems. There is much literature noting the benefits of large displays that helps justify this focus; for example [5–7].)

Someday, high-resolution 3 m-by-3 m 4K-resolution panel displays formed into a ceiling, positioned vertically as walls, and tough enough to walk on, may allow CAVEs to be built in just about any enclosed workspace, given suitable ventilation and provision for safe ingress and egress. Today, however, a 6-sided cubic-format CAVE presents many challenges to a panel-based

implementation. As a result, less-than-totally surrounding panel-based systems have been developed and installed by the authors to offer significant and usable VR imagery, given the current technological limitations. Someday a perfect technology may emerge, such as seamless, lightweight, high-resolution fabric or painted-on wall, floor, and ceiling coverings. Of course, software for control and content will always be a challenge, and that is worthwhile for researchers and developers to address while anticipating hardware advances.

The CAVE was envisioned from the outset as a device to enable distance collaboration among viewers immersed in shared computer-generated scenes – a kind of 3D telephone booth [8–18], a technique called *tele-immersion*. Much work has gone into building and optimizing ultra-high-speed computer networks suitable for sharing Gigabits/second (Gb/s) of information across a city, region, nation, and indeed, the world [19–26]. Tele-immersion development and application is facilitated by a 10Gb/s network link between KAUST, EVL, and Calit2 via Amsterdam as an extension of the Global Lambda Integrated Facility [27].

1.2. CAVE Goals and Features

The ideal VR viewing device has been visually portrayed in movies, from the iconic first *Star Wars* movie from 1977, and the Holodeck in *Star Trek*, and arguably before. Such systems usually include augmented (see-through) reality as well so that the 3D immersive correct-perspective imagery is seamlessly merged into the surroundings, when desirable. Of course, all of the visuals in a movie are computer-generated special effects, shot or generated from the movie camera's perspective, and not the viewer's (or actor's) perspective. Movies are projected on a flat screen, so the portrayal of a movie's "VR" display is easily faked. Figure 2, a composite concept, shows the goals of a fully surround autostereo VR system that does not yet exist.

The ultimate hardware design goals for a VR system are:

- Compact footprint, so it can fit in work or home spaces
- Scalability, so it can be laptop, desktop, corner-sized, room-sized, or stadium-sized
- Usability in normal room lighting, so that other work/play can go on simultaneously
- Low noise signature, so that people can talk and generated audio can be heard
- Low thermal signature, to minimize need for ventilation and cooling




Figure 1. The “classic” CAVE is a multi-person, room-sized, high-resolution, 3D video and audio environment. Early CAVEs had less than a megapixel per wall per eye; newer ones have up to 15 times as many pixels. Photo: National Center for Supercomputing Applications at University of Illinois at Urbana-Champaign. Note: All the photographs, unless otherwise noted, in this paper were shot with the stereo turned off so the images are not doubled.




Figure 2. A composite rendering of the EVL/Calit2 vision of a future collaboration room. Illustration: Jason Leigh.

- Potential to hold several users, and/or be network connected, to allow collaboration
- Extended service intervals and easy access for maintenance, to reduce expense
- Power-efficient, to reduce cost and cooling
- Articulated, easily shippable screens and rapid installation/de-installation for field and traveling exhibit use
- Low cost, so that ubiquity is conceivable.

Desirable features include:

- High resolution, so virtual images are seen as sharply and in as much detail as in reality
- High brightness and contrast, so colors are vibrant and not washed out or dim
- Production of computer graphics and the display of captured imagery that is equivalent to or exceeds human visual acuity, in 3D with the correct viewer-centric perspective rendering for every viewer
- Input and full recognition of the viewer's or viewers' being and actions, including speech, non-verbal utterances/noise making, and gestures

If useful to the task,

- Audio (sonification) at or exceeding human aural acuity, fully surround, listener centered and focused
- Touch (tactile) input from the user and touch output from the VR system, allowing haptic input and feedback, for all users, for example [28]
- Olfactory (smell) output delivered to each user, and input recognition as well
- Taste output and input recognition

Linking such devices together with near-zero latency and no noticeable compression artifacts is a goal, as is providing these affordances with no user encumbrances (special glasses, headphones, nose tubes) except as obviously required for touch and taste.

Much as color television replaced black and white television, VR practitioners want to reduce the objections to VR down to the point that 2D visualization devices are no longer for sale; as with color, one buys a display and gets 3D for free, but can turn it off when preferred. To do this, displays need to (at least appear to) be large (that is, the edges are generally outside the

viewers field of view) and the viewers need to be tracked accurately. Such VR capability is becoming available, even in the home, although the price is not yet asymptotically approaching the cost of 2D displays. Audio capability, however, is somewhat more available to consumers, and tactile input/feedback has long been a feature of consumer gaming controllers of various sorts. Taste and smell are still challenges, occasionally investigated. Tele-immersion (linking VR devices over networks) is routine now with latencies down to nearly what is dictated by the speed of light in fiber, although compression is still needed since it is not really possible to transmit a CAVE's worth of bitmaps as uncompressed video. Typically, for tele-immersion sessions, most of the geometry and texture maps are downloaded ahead of time, and some video, audio, and interaction information is sent in real-time.

2. Beyond Classic CAVEs

Improvements to the Classic CAVE have mainly been the result of commercialization, better projectors, graphics, and software. Two CAVE-like systems have been built that greatly improve the resolution by tiling and using higher-resolution projectors: the StarCAVE and the Cornea/C-6.

2.1. StarCAVE

The StarCAVE [29] (Figures 3, 4), designed and built at UCSD/Calit2, is a 5-wall plus floor projected VR room, operating at a combined resolution of ~68 million pixels, ~34 million pixels per eye, distributed over 15 rear-projected wall screens and 2 down-projected floor screens. The StarCAVE is a surround system with diameter of 3 m and height of 3.5 m. Its 15 wall screens are acceptably non-depolarizing, high-contrast, low hot-spotting rear-projection screens, stacked three high, with the bottom and top trapezoidal screens tilted inward by 15 degrees to increase stereo separation, important for a good VR experience. The screens are 1.19 m high by 2.13 m wide, narrowing to 1.66 m at the top and bottom. The pentagonal floor images are projected from overhead onto a non-depolarizing, coated floor. User interaction is provided via a wand and multi-camera, wireless tracking system¹.

¹ Note: the tilting of the screens allows more on-axis viewing by a viewer standing at the center than if the screens were purely vertical. Rear-projection




Figure 3. Calit2's StarCAVE from above. Photo: Tom DeFanti.




Figure 4. Calit2's StarCAVE. Illustration: Greg Dawe.



Figure 5. Cornea with door screen open. Content: Tom Levy, Kyle Knabb, and Jurgen Schulze. Photo: Tom DeFanti.

screens tend to depolarize when viewed severely off-axis—vertically or horizontally—so tilting helps minimize vertical depolarization, thereby helping to minimize ghosting. Depolarizing is lessened by the larger than 90° angle between screens at the corner/corner and corner/floor intersections in this pentagonal cross-section, tilted-screen CAVE. Tilting the top and bottom screens also helps make the illumination to the screens more uniform appearing because the projectors are closer



Figure 6. Cornea with users. Content: Tom Levy, Kyle Knabb, and Jurgen Schulze. Photo: Tom DeFanti.

2.2. Cornea

Cornea (Figures 5–8) is a Mechdyne-designed and built 6-sided ~ 16 -Megapixel/screen CAVE-like large-space VR system built first at Iowa State University as the C-6. C-6/Cornea's $\sim 4,000 \times \sim 4000$ projected pixel resolution per screen is a significant improvement over the classic CAVE resolution of $\sim 1000 \times \sim 1000$ per screen.

Design objectives for the Cornea were to:

- Create the highest resolution, brightest CAVE possible in 2009
- Fully enclose the participants in visuals
- Provide high-fidelity user tracking
- Incorporate spatialized sound including variable, interactive model-based reverberation
- Allow tele-immersion via 1-10Gb/s networking
- Capture and transmit sessions with switched and transmissible audio and HD video

Cornea is currently, in 2010, the world's brightest virtual environment with the highest resolution (The Iowa C-6 is the same resolution but half the brightness). Each of Cornea's six ~ 3 m-by-3 m screens displays

to the optimal on-axis position to the viewer than if they the wall panels were strictly vertical. This is quite measurable: the screens used in the StarCAVE were custom developed and subsequently chosen for acceptably low hot-spotting as well as low ghosting by relying on quantitative measurements which turned out best when the spot meter was on-axis; additional details and benefits of this screen configuration and material are found in [29].

~15-Megapixels/eye (there is overlap for blending between top and bottom projector sets, so each screen is not quite twice 2K lines). This amounts to a total resolution of ~90-Megapixels/eye (of course, since one cannot see behind one's head, at *least* half of the 90-Megapixels/eye is unseen by any viewer at any time). 24 4K 10,000-lumen Sony SXR [30] projectors (4096 x 2160 pixels each, not including blending overlap) are used, specifically, 4 SXRDS illuminating each of 6 screens, including the ceiling, floor, and a closable "rear" wall screen (these projectors have an aspect ratio of 2:1, so 4 of them gives a square wall image in stereo, see Figure 7). A 24-PC cluster is coordinated by a head node. Each of the 24 PCs has a four-output outboard Nvidia Quadroplex, thus 96 5600-model synchronized ("g-sync'd") GPUs are used to drive the displays.

Calit2 audio researchers designed the advanced sound system for this a six-sided virtual environment, directing Meyer Sound Lab engineers to neither obscure the visualization screens with speakers, nor place the speakers too far away, since VR screens can be somewhat acoustically opaque. To contend with these challenges, the Cornea is equipped with 28 speakers and 4 subwoofers outside the screens. The sound system has a sampling rate of up to 96 kHz, which is more than two times the sampling quality of a compact disc (44.1 kHz). A higher sampling rate means the system offers more flexibility in terms of audio playback, and high definition performance consistent with state-of-the-art audio production techniques and equipment. Papers describing this audio system are in review and should be easy to search for once published. Meyer sound has created a brochure about the KAUST installation [31].

Cornea makes it possible for users to record their VR experiences, both visually and aurally. The facility can stream 32 channels of high-definition audio and video from Cornea to KAUST's Interactive Media Room, where it can be recorded and archived. This capability provides a needed tool for scientists to collaborate with their counterparts at institutes that lack virtual environments. Mechdyne engineers designed and implemented this recording capability.

The Cornea's audio system is driven by custom audio switching, routing and rendering software that incorporates the latest advances in realistic virtual acoustic imaging, such as air absorption, Doppler effect, source directivity, reflection/absorption simulation, and convolution reverberation. This system functions both as an audio rendering engine and as a creative sound-design authoring tool for Cornea visual applications. KAUST scientists can create audio soundtracks that correspond to their 3D visual data, so a walk through a virtual

architectural rendering, for example, can include the sound of one's own voice and footsteps reverberating off the virtual walls.

KAUST has two additional SXR-based VR facilities for virtual reality research, development, and application use. One is a 75-seat room with a 32-Megapixel/eye stereo VR screen (precisely two walls of the Cornea side-by-side, involving eight 4K projectors in a 2x2-paired layout plus a ninth projector in the middle). This multi-purpose room also features an advanced audio research capability whose description is outside the scope of this paper. The second one is a development system, the DS-1, which is one-half wall of the Cornea, used primarily by programmers.




Figure 7. KAUST's Cornea. Illustration: Mechdyne Corporation.



Figure 8. Cornea upper projectors. Photo: Greg Wickham.

3. Design Challenges of CAVEs, Present and Future

3.1. Projector-based CAVE Design Challenges

Projection-based CAVEs have many design constraints that prevent designers from approaching the goals specified in Section 1.2, and thus limit CAVE ubiquity, namely:

- 10 m-by-10 m-by-10 m or larger spaces are not commonly found or easy to justify; buildings normally need to be built around such spaces at great expense. CAVEs are architectural achievements as well as technical installations.
- Power and air conditioning needs of projected CAVEs are high (e.g., 25-100 kW), although tiled display walls of similar size can consume as much.
- CAVE projectors are expensive, heavy, and, as with all projectors, require re-lamping every 800-2000 hours of use. Lamps range from several hundred to several thousand dollars each; projection-based CAVEs have 4-24 projectors with one or two lamps each.
- CAVE screens face each other and thus, unlike single screens in movie theaters, pick up light from each other and are further illuminated by any ambient lighting, which means they are best viewed when the ambient lighting is very low and the screens themselves are made of dark-surface projection material. Projectors also have hot spots and off-axis intensity attenuation issues, which can be compensated for with extra rendering, given that the location of the tracked viewer is known [32].
- Hard-surface rear projection screen systems have odd sounding acoustical signatures due to uncontrolled reflections, standing waves and room modes. They also inhibit the ability to project sound directly to the listener, meaning that speaker-produced sound will tend to reinforce the room's odd acoustical quality. Softer surface rear projection screens are better sounding; they can pass some sound directly to the listener. In both cases, significant analysis and equipment investment is required to address the screen-based acoustical anomalies.
- Alignment of the projectors is difficult and time consuming, more so as the number of projectors increases and one wraps around in all

dimensions. Cheaper projectors typically have no auto-alignment features to help this, and are built into cases that warp a bit as they warm up, further frustrating the goal of pixel-perfect alignment. Expensive projectors with sophisticated mounting and auto alignment systems are better, but the differential expense per projector (perhaps 5-10x) has to be multiplied by the number of projectors to appreciate the additional cost. Software can mitigate the problems if performance penalties incurred by the extra rendering/compositing needed are acceptable. This is an area of continuing research.

A primary goal for the CAVE of the future is to allow CAVE-like surround VR in a normal work or home environment, that is, in a ~3 m-by-3 m-by-3 m or even smaller room (most office/lab spaces have nearly 4m between the floors, although 1 m of that is often taken by ducts, pipes, and conduit). In order to do that, projection needs to be eliminated, and the CAVE of the future has to be constructed from self-illuminating panels of some sort.

A further goal of the CAVE of the future is to offer contrast (and saturation) as good as modern home HDTV displays, that is, greater than 1000:1. Internal light spillage and reflections attenuate contrast in projector-based CAVEs due to the screen material typically used. Bright surround scenes offer grays at best, not blacks, and rather unsaturated (pastel) colors. Consumer panel displays, it should be noted, have coatings that are effective in reducing the effects of spillage and reflection.

3.2. Panel-based CAVE Design Challenges

Panel-based CAVEs have the following design challenges:

- A ~3 m-by-3 m-by-3 m display would need to be built using many panels, given current (2010) display size and resolution. Consumer 3D-capable panels, however, are modestly priced (\$2,000-\$5,000 depending on size), so cost isn't the critical issue (compared to the expense of classic CAVEs).
- 3D-capable consumer panels are designed to be viewed conventionally, that is, not severely off-axis in the vertical domain [33]. This means that a more faceted tessellation of the space with screens, as in the StarCAVE is useful so that the screens are arranged more perpendicular to the viewer's line of sight. However, since panels are only produced in rectangular formats, non-cubic tessellations have to

incorporate overlapping panels, as in the NexCAVE (Section 4.1.1).

- Consumer panels typically have large (>35 mm) bezels around the display, and bigger ones on the bottom. This produces a window mullion-like effect when they are tiled. Commercial ultra-narrow bezel displays minimize this problem greatly.
- Ceiling and floors are an issue. Panels are made of glass and are 20-50 kg so an infallible structure would need to be provided to securely hold panels over the heads of users. Floor displays could be viewed through a clear acrylic screen thick enough to support people. Non-cubic tessellations of panels make the floor and ceiling a design challenge, although not an inconceivably difficult one.
- Directional audio in a highly polished, acoustically reflective glass/acrylic enclosed space is hard to do. Headphones/earbuds for the viewers, rather than speakers, are a solution, albeit an encumbering one.
- Access must be considered. Projection CAVES, such as Cornea and StarCAVE, which totally surround viewers, have movable screens for one wall. One column of panels could be hinged and made movable to allow human access, effectively creating a door, albeit a heavy one.
- Airflow must be considered as well. Any fully enclosed CAVE, including Cornea, needs either a ventilation system or a door opened after a short period of use (~15 min) to let in fresh air. Panels project significant and palpable amounts of heat into the space (not an issue with projectors which are outside the viewing space), so a source of cooled airflow or at least good convection would be needed for constant use.

Building a room-sized VR display with panels clearly has its challenges, some of which can only be met with compromises. Section 4 describes those compromises that the authors have chosen to implement to date, with the intention that some of these will be ameliorated by future display technology in mass production, thus facilitating more wide-spread design and installation of fully-surround VR systems.

3.3. Active Stereo and Passive Stereo using Special Eyewear

3D movie theater projector and consumer 3D HDTV manufacturers provide competing passive and active stereo

implementations. In 2010, polarization does not seem to be the choice of the majority of manufacturers for home 3D HDTVs (both plasma- and LCD-based). The consumer market is dominated by field-sequential (alternate-field) “active” stereo TVs. Active stereo glasses [34] are made of LCD shutter material that switches clear/opaque in synchrony with the projector or 3D panel left/right images. The glasses are bulky (although getting lighter, cheaper and more fashionable as consumer versions appear), require frequent battery replacement or recharge, and the synchronization sensor must be visible in line with its source. Other constraints that inconvenience or defeat use in future CAVES are only now being discovered. For example, the new (May 2010) 55” Samsung consumer 3DTV uses linearly polarized active glasses that make the display go black as one’s head is tilted, so these display panels cannot be vertically oriented (i.e., “portrait” mode) to make a multi-panel surround CAVE system (unless the glasses polarization is circularly shifted), nor can they be used in floors or ceilings.

When developing multi-tile systems, all the panels must be synchronized so that every screen has the same left/right eye as the others in view. Users in living rooms or laboratories may have an incompatible array of 3D screens in their multitasking lives in the foreseeable future; 3DTVs, mobile phones, clocks, smart appliances, and laptops cannot be easily synched until a universal clock is developed, similar to “house sync” in a TV studio, for all such displays in an environment. Today, one cannot assume one manufacturer’s 3DTV’s active glasses will function with another’s.

Of course, passive polarization only simplifies the problem by eliminating the need for left/right eye shutter synch—each brand and model of 3D LCD display typically has different polarization angles in its design, making multi-tasking difficult until a single standard for polarization is adopted as well. Even in a fully-surround VR environment, multi-tasking users will want to use their legacy 2D and 3D mobile phones, electronic notepads, GPS navigators, and LCD wristwatches, so the problem will linger.

3.4. Autostereo: Stereo without Special Eyewear

The ideal VR display would function totally without viewer encumbrances. Without special VR glasses, there are no synchronization or polarization issues to prevent users from multi-tasking; for example, viewers could see 3D while taking notes on paper or looking at other computer screens. EVL has extensively researched autostereoscopic display systems for over twenty years [35] (Figure 19)

and has prototyped static and dynamic parallax barrier systems (Varrier [36, 37], Dynallax [38–41]). The primary benefit of an autostereoscopic system is that it eliminates the need to wear special 3D glasses, which is associated with all other types of VR systems. Current systems, however, make some assumptions as to where the viewer is, which frustrates the idea of a multiplicity of autostereo displays being useful. Tracking, that is, being able to tell the display hardware/software where the user is, can help.

3.5. Visual Acuity in CAVes

Classic CAVes which typically have a ~ 1 -megapixel single projector covering each 3 m-by-3 m screen, provide a rather low acuity image to the viewer in the normal viewing position (the “sweet spot”) due to the 3 mm pixel size. Using a standard vision-testing chart placed on the screen, three of the authors with slightly better than normal vision noted which lines on the eye chart they could see from the sweet spot on the VR displays discussed in this paper, three of which are described above (Cornea, StarCAVE, CAVE), and two below (NexCAVE and REVE). Using 1 arc minute per pixel, the acuity was also calculated for each display and a classic CAVE using Christie DLP 1280x1024 projectors (of which only 1024x1024 is used if the screens are square).

Table 1. Visual Acuity.

	REVE	Cornea	NexCAVE	StarCAVE	CAVE
Observed	20/16	20/32	20/24	20/40	–
Calculated	20/16	20/34	20/24	20/48	20/137
Sweet Spot Distance	3 m	1.5 m	1.5 m	1.5 m	1.5 m

Note that in this terminology, normal vision is 20/20 (6/6 in metric) and legally blind is 20/200, which refers to the size of letters one can see from 20’ away on a Snellen eye chart (please excuse the short lapse into US measurement for this section!). The numbers imply that at 20 feet, one can see what a “normally-sighted” person can see at 20’, or 200’, respectively. (“20/20 vision is in common use in the US; Wikipedia, which has a rather complete discussion, notes: “in some countries, acuity is expressed as a vulgar fraction, and in some as a decimal number.” [42]). Since these VR systems are much smaller than 20’, we have scaled the results to fit the sweet spot working distance. Note that if one stands ~ 3 m away from a screen, that is, with one’s back against a wall, instead of ~ 1.5 m, in the center, the acuity of the Cornea, NexCAVE and StarCAVE is equal to or better than 20/20, which is pleasingly non-fuzzy to normally-sighted viewers. The

24” Alioscopy display gives excellent projected acuity at 3 m, an interesting result.

Also note that using 4K (4096x2160) panels/projectors for the NexCAVE/StarCAVE screen sizes instead of 1920x1080 HDTV panels/projectors would give $\sim 20/20$ acuity at the 1.5 m sweet spot, an argument for 4K technology in this application, if one can be convinced that it is worth the order of magnitude increase in cost and complexity, given the price of currently marketed 4K panel/projector technology and the need for multiple synchronized GPU cards for each 4K device. (Of course, bright 1280x1024 displays were an extravagant expense a decade ago, as were the Silicon Graphics, Inc. computers that drove them). If narrow bezel 4-megapixel active or passive stereo panels or projectors were mass produced as are the popular 4-megapixel 2560x1440 30” Apple Cinema Display and its Dell equivalent, they would be quickly adopted because they would be inexpensive and modern game graphics cards provide 2560x1440 resolution output.

The point here is to attempt to put some quantification into what one perceives, especially with static imagery in these VR systems, that is, when the motion stops; it really makes a qualitative difference when the screen images are seen as perfectly crisp, that is, with a perceived 20/20 acuity. Once the images start moving, resolution becomes less important, and reading text or numbers normally becomes irrelevant, which is why VR systems with 20/30 or 20/40 or even 20/137 visual acuity have been compelling enough to build. Nevertheless, when the user stops navigating, and in some cases, the detail fills in further by refinement, it is distinctly visually rewarding if the image is seen at the fullest resolution possible. (Whether a CAVE with 20/100 acuity is 5 times worse, subjectively, than a CAVE with 20/20 acuity is beyond the scope of this paper, but a worthy topic to consider and formally verify through a rigorous study with human subjects.) In addition, for further inquiry, it would be useful to actually place the eye chart virtually at 20’, not on the screens, since this would also test the impact of tracking and projector alignment.

4. CAVes Constructed with Panels

If one accepts that a future of the CAVE involves moving away from projection-based systems, flat panels are the realizable first step. To create large-space panel-based VR systems in 2010, one must effectively tile panels and have them display left and right eye images in usable ways. Tiled mono displays are now common [43] (Figures 9, 10), but stereo-capable panel displays have

just become available in 2009.

This section discusses implementation issues with panel-based tiled displays. No VR can be seen without working hardware. Several built and one system in the design phase are described below which incorporate and exploit panel display technology advances. Not one of these panels was initially designed by the manufacturers to be used for stereo tiled displays. (Projection-based bezel-free tiled displays have been around for 20 years, of course, and arrays of standard-definition TVs have a long history.)

The frames (called bezels) that manufacturers put around panels for electronic, structural, esthetic, and other design reasons create obvious seams in the imagery when panels are used in arrays. Narrow-bezel 720p mono displays are now available and are discussed in Section 4.2. The (nearly) seamless 3D tiled display panel is coming, but no one knows when.




Figure 9. An OptiPortal tiled display showing EVL SAGE images at KAUST coming over a 10 Gb/s network from Chicago. The OptiPortal [43] is a tiled display that is the common visual interface to the OptiPuter global-scale computing system [44]. OptiPortals are designed to allow collaborative sharing of extremely high-resolution graphic output, as well as video streams over 1-10 Gb/s networks. OptiPortals typically consist of an array of multiple LCD display panels (1-to-4-megapixel each), driven by an appropriately sized cluster of PCs with optimized graphics processors and network interface cards. Photo: Tom DeFanti.



Figure 10. Conventional thick-bezel OptiPortal mullions obscuring the image. The CGLX VideoBlaster module is used. Video content: April Bailey. Photo: Tom DeFanti.

4.1. Conventional-Bezel Panel-based CAVEs

CAVE-sized stereo can be achieved using panels in an array in a variety of ways as noted above: passive stereo, active stereo, and autostereo. (One can easily achieve stereo with two displays with a half-silvered mirror between them, but this technique is not scalable to an array of displays). The passive stereo panel-based NexCAVE and two autostereo panel-based systems are discussed in this section.

Using active stereo panels for more than 3 panels is an area of great interest since the consumer stereo active HDTVs are the most inexpensive and currently mass-produced solution; however, they lack obvious means of synchronization of left/right eye, a necessary feature. An active stereo, single panel tracked VR display avoids this problem, of course [45, 46]. Once the means to synchronize active stereo left/right eye images is provided by the manufacturers or discovered by hacking, active stereo panels will be quickly built into panel-based CAVEs. (Note that Classic CAVEs are built with active stereo projection. Synchronization of left/right eye images was a feature of the Silicon Graphics, Inc. Onyx multi-processor, multi-graphics engine systems that provided 96-120 frame-per-second stereo; the feature was specifically added for the CAVE. The Cornea is run in active stereo mode, achieved by large synchronizing shutters placed in front of the SXRD projector lenses. Some projection VR systems have also used passive stereo.)

4.1.1. NexCAVE: A Passive Stereo Panel System

KAUST currently has a 21-panel NexCAVE, (which stands for "NewXpolCAVE," with "Xpol" being short for "micropolarization") passive-stereo 3D environment (Figures 11-15). Micropolarization ("Xpol") is a technique to create stereo images by changing the polarization of the video image on a line-by-line basis, alternating between right and left circular polarization [47-49]. A 1080p HDTV can have affixed to its surface a transparent overlay of 1080 lines, 540 polarized one direction interlaced with 540 polarized the other direction. Besides the relative simplicity of manufacture, it is passive and uses circularly polarized glasses. Also, it is a spatial technique, so that there is no need to synchronize left and right eye images, since they are merged (unlike field-sequential stereo, which requires active stereo glasses and, if there are more screens, universal left/right synchronization.) For this reason, Xpol displays are scalable using the same screen-to-screen synchronization techniques as any tiled display.

The availability of commercially produced Xpol LCD displays is a recent (2009) advance in marketing consumer



Figure 11. Tom Levy and Sami Maghlooth in the KAUST NexCAVE reviewing the VR reconstruction of UCSD's excavation at Khirbat en-Nahas in southern Jordan. Content: Kyle Knabb and Jurgen Schulze. Photo: Tom DeFanti.




Figure 12. KAUST's NexCAVE showing a stereo 360° scan of the Wisconsin State Capitol (in mono mode). Content: Dick Ainsworth. Photo: Tom DeFanti.

HDTVs. The JVC Xpol 46" display allows 3D generated or scanned stereo images to be produced in real time as well as played back as HDTV streams. The NexCAVE at KAUST consists of 21-tilted displays (arranged in a 3-high by 7-wide configuration (Figure 12)), with the top and bottom tiles tilted inward in order to help preserve stereo (since these displays have a limited vertical angle of view). GPU-hardware accelerated stencil buffer routines are used to combine left- and right-eye views into a line-by-line alternating image prior to display, at real-time frame rates. This operation is similar to that performed during real-time autostereo interleaving [50].

Calit2's NexCAVE has 10 panels (Figure 13) (a 3x3 array with an additional display at the bottom of the middle column) and was designed for shipping so that it can be deployed in research exhibit booths at conferences. Because the NexCAVE operates well in ambient lighting and is free standing, installation/maintenance in a booth or any other work/play space is greatly simplified. Of course, any complex display like this in a museum or unguarded public setting would need further enclosure and protection from tampering.



Figure 13. The 10-panel NexCAVE at the SC'09 conference showing content developed by Philip Weber. Photo: Tom DeFanti.

The tilted StarCAVE non-rectangular screen tiling helped inspire the design of the NexCAVE. Polarized screens tend to ghost more (that is, the stereo separation attenuates) when viewed at an angle. The tiles need to be arranged to allow approximately on-axis (that is 90° perpendicular to the screen) viewing as much as possible. The JVC Xpol display has a broad off-axis horizontal stereoscopic viewing angle of about 140°, but a vertical viewing angle of only 20°. The StarCAVE's upper and lower sets of screens are actually cut in the shape of trapezoids and tilted in, so that the viewer's line of sight is fairly perpendicular to the screen in typical viewing positions, and the design of the framing makes the unlit bezel in between screens negligible [29]. Trapezoidal LCD panels are conceivable, but not likely to hit the market anytime soon, so the NexCAVE's required tilted-in screens are physically achieved by overlapping JVC Xpol screens, including positioning the bezels behind one another, which helps minimize their perceived thickness (Figures 14, 15). This arrangement works well and is not disruptive when viewing 3D scenes if one does not specifically pay attention to the bezels. The NexCAVE's overlapping panels are not very visually distracting in practice. (A different tiling is proposed for the NG-CAVE, Section 4.2.2). When some detail is blocked by a bezel, the tracked viewer just naturally moves a bit to look around the bezel, just as one would look out a window with mullions. Also, just as in the case of a window, vertical bezels are perceived less than horizontal bezels for virtual objects which do not happen to be in the screen plane, because due to the horizontal offset of

human eyes, at least one eye can usually see the object at all times. (All accommodation information in any VR system is essentially incorrect, unless the distance of the viewed object happens to be at the same distance from the viewer as the screen, so having a fragmented image wrong in a variety of places instead of one place is not any more incorrect or disturbing.) Of course, narrower bezels would be better, as in the AESOP display and NG-CAVE discussed below in Section 4.2.

So far, a NexCAVE with a full floor and ceiling has not been attempted, although with the appropriate structure for the overhead tiles and a strong and clear surface to stand on, a more surrounding NexCAVE could be built. The space between the overlapping NexCAVE screens provides an opportunity for forced air ventilation, eliminating one design/usage problem with fully surround CAVEs mentioned above.



Figure 14. The back side of the KAUST NexCAVE. Photo: Tom DeFanti.



Figure 15. The KAUST NexCAVE turned off to show the structure designed and built by Greg Dawe. Photo: Tom DeFanti.

The NexCAVE advances VR in three important ways over projector-based systems. First, as noted above, dozens of projectors are really hard to accurately align and keep aligned due to the dimensional instability of the projectors, supports, and screen systems. The result is that the user's eye/brain system needs to do the "final" alignment. Since the NexCAVE panels' left- and right-eye images are perfectly aligned, the eye/brain fatigue caused by imperfect alignment in projected classic CAVEs and the StarCAVE seems to be greatly mitigated according to users. Ideally, an experiment by independent researchers needs to be run to verify this opinion.

Second, the contrast offered by the NexCAVE has been measured by the first author to be in excess of 300:1, or 10 times that of the StarCAVE (30:1) and 100 times that of the Cornea (3:1) (Figures 16-18). The measurements in each case were taken using a Panasonic Lumix LX-3's internal light meter with the lens wide open, noting the shutter speed readouts, an arguably less precise instrument than the professional spot meter used in the StarCAVE screen development mentioned above, but surely sufficient to report order of magnitude differences. The whites should be as white as a white shirt, and the blacks as black as a black pair of pants. (The lack of contrast in classic CAVE screens indeed motivated the development of new screens for the StarCAVE as described above and in [29], but it would be difficult/expensive to make rigid screens large enough and to suspend them overhead to create the classic CAVE's 3 m-by-3 m walls. (The StarCAVE's ceiling is open.) Experiments are underway to extend the size of the StarCAVE screen technology to 3 m x 3 m using a semi-rigid material that can be rolled for shipping.)

Third, the articulation and modularity of the screens in the NexCAVE are a contribution to the future of the CAVE, in that shipping, installation, expansion, and de-installation are rather straightforward compared to classic CAVEs.

4.1.2. The Varrier Autostereo System

The Varrier (Figure 19) is a head-tracked, single-viewer, VR system that generates high-quality images with 5% ghosting [37, 51], and supports a depth range that goes from 1/3 meter to infinity. It does, however, have several drawbacks. Varrier is a single-viewer system, so other viewers in the space see a very degraded image, which is disturbing. Also, Varrier requires a high frame rate, low-latency, high-accuracy tracker without which viewers see pseudo-scopic or darkened images. As currently implemented, Varrier is suitable for a single viewer who is either seated or not moving too quickly within the space (as the computers need to continually update the graphics



Figure 16. KAUST Cornea checkerboard pattern. Photo: Dan Sandin.




Figure 17. Calit2 StarCAVE checkerboard pattern. Photo: Jurgen Schulze.

from the user's perspective, so the computations need to keep up with the viewer's movements!). One reason the CAVE was replicated was that it nicely supports small group viewing and has much less sensitivity to tracking, especially compared to head-mounted displays and Varrier. As Varrier does not support multiple users, it might be best to further develop it as a personal VR system, perhaps integrated into a semi-cylindrical-format 2.5 m-diameter workspace.

4.1.3. The REVE Autostereo System

Calit2/EVL is currently experimenting with multi-viewer autostereoscopic systems that use lenticular parallax panoramagram (LPP) technology, such as the Alioscopy display, a commercial product. LPP has the potential to reduce sensitivity to tracking, as it splays out a number of views (8 for Alioscopy) so that in a head-tracked configuration, the viewer is always near the center of the fan of images. If the viewer moves quickly, he/she moves into adjacent views that are still in correct stereo. The more modest tracking requirement, if used, means that a camera-based facial recognition system is a rational




Figure 18. KAUST NexCAVE checkerboard pattern. Photo: Daniel Acevedo.

choice; the viewer would not have to wear targets, and would still see the correct VR perspective. A version of Varrier that used facial recognition tracking with success is described in [52].

Since LPP views are multiplexed out into space, the 3D experience for non-tracked viewers is of acceptable quality, unlike the Varrier. The major limitation of the current state-of-the-art multi-viewer head-tracked and non-head-tracked LPP displays is the limited depth,




Figure 19. 60-panel Varrier autostereo display at Calit2. Mars data interactive rendering: Bob Kooima. Photo: Tom DeFanti.

which is approximately ± 1 m. The depth is limited because engineers had to make tradeoffs in the design. The small number of angular samples through the 3D space to be displayed (8 for Alioscopy) produces angular aliasing; the aliases are fuzzed out using an optical reconstruction filter that uses ghosting (cross-talk) to make highly aliased images look better [53]. However, this ghosting is significant, that is, greater than 20% (Figures 20, 21). If an object has a high disparity, the viewer sees multiple images.

Instead, a different set of tradeoffs can be used to design LPP displays. Specifically, more samples can be taken in the angular dimension, thereby making the increments in the angular samples smaller. Then, a proper reconstruction filter can be implemented in software [54]. While this method makes the spatial resolution lower, one can compensate by selecting displays with higher pixel pitch and/or by tiling the displays to get higher resolution. This gets around the limited depth of the current generation of LPP displays, and should allow researchers to significantly expand the range of depth that these displays can handle.

VR encumbrances, like special glasses, would be best eliminated in that synchronization and polarization also inhibit ubiquity. Scalability is also desirable, as multi-tile displays are the only way to achieve resolutions exceeding 4096x2160 pixels with technology in the near future. To explore scalable autostereo using commercially available displays, Calit2 designed an autostereo display called REVE (Rapidly Expandable Virtual Environment) (Figure 22 [55]) in 2009. (It is “rapidly expandable” in the sense that one can, with physical installation and expense, but no re-design, make it larger quickly. “Reve” is French for “dream” or “day dream”, recognition by the first author that the inventor and his company are French.)




Figure 20. Too much z-depth in the model causes ghosting. Content from the “How Much Information” project, Andrew Prudhomme. Photo: Tom DeFanti. (Autostereo not turned off.)




Figure 21. Compressing the depth by 6:1 eliminates the ghosting. Content from the “How Much Information” project, Andrew Prudhomme. Photo: Tom DeFanti. (Autostereo not turned off.)

REVE is a flat wall of multiple 3D autostereo LCD displays based on LPP technology. The KAUST REVE has 18 42” panels in a 6x3 array, making a display 6.11 m-by-1.83 m in size. Several competing lenticular systems were evaluated, and then Calit2 prototyped tiled configurations of three, six, and nine 24” Alioscopy displays [56]. In addition, new display drivers were written, software adapted, and techniques for synchronized playback of multi-screen, pre-computed, compressed animations were developed. This was effectively an all new effort, building on the past 20 years of autostereo research.

REVE’s software is interactive, making it possible for viewers to control parameters directly in real-time scientific computations with precision and high resolution, at a speed of 30+ frames per second. Consumers need this capability for 3D games, among other uses. As with any spatially multiplexed display, and especially with LPP ones, resolution is a fraction of HDTV; however, resolution is regained in REVE, obviously, by tiling the displays.

A current shortcoming with the REVE display is its narrow z-depth, which limits its use to shallow 3D objects like friezes and etchings; typical VR fly-through images are not successful. Calit2/EVL developed a real-time depth compression technique that



Figure 22. The KAUST REVE 18-panel display using CGLX-based CineBlaster3D video software. Julia4D Animation: Dan Sandin. Photo: Tom DeFanti. (Autostereo not turned off.)

automatically compresses the 3D scene to the usable z-axis range, which works well with some types of images, but not with others (Figures 21, 22). This issue is made worse by tiling, not better, unfortunately, since even as the image's x and y dimensions increase, the useful depth does not, making the resulting images more comparatively compressed in z, and thus more flat-looking. Nevertheless, the images are more dimensional than 2D representations of the same data, if carefully designed for the z-axis constraints.

Another issue with the REVE display is that the screens present autostereo in a cyclic way; in the case of the Alioscopy display, it is 8 views that repeat (Figure 23). This design feature means that many people can easily see a good 3D image from multiple optimal viewing positions, arrayed in repeating fans about 2/3 meter wide. If a viewer moves such that one eye is in one fan and the other eye in the next fan, a vertical fuzzy bar appears somewhere in the scene. To the right and left of the bar, stereo is still correct. If the viewer moves 10cm or so the bar moves out of the scene. Tracking the user accurately, as in the Varrier display, can eliminate the bar, and only slightly degrade the experience for the other viewers. Tracking requires markers on the participants in the current implementation, which is an encumbrance. Fully marker-less tracking [52, 57] is desirable to integrate into future CAVEs.

A fully or even partially surround CAVE would be hard to implement with the current Alioscopy display and provide meaningful VR. For example, the narrow depth combined with a cylindrical configuration would allow wrap-around friezes, but not deep 3D scenes.

Alignment of the images in the REVE's multiple tiles was achieved in the same manner as for the Varrier [37]. Synchronization of the images is simple since stereo, as in the NexCAVE's stereo, is spatially multiplexed, unlike temporally multiplexed screens that need to keep the




Figure 23. Alioscopy technology is based on 8 distinct points of view (POV) each of which is aligned to the LCD's sub-pixel. 2 of the 8 points of view can be seen by the user in a sweet spot. Each viewing angle is 8 degrees from sweet spot to sweet spot. Illustration: Alioscopy USA, Inc.

right- and left-eye data perfectly in synch, as discussed above.

Special 3D glasses are not required to view data in three dimensions, so the display is an excellent way to show carefully designed shallow-depth data, in motion or still, for groups of people, especially casual passers-by. However, severe off-axis viewing is problematic (Figure 24).

Video playback is possible on tiled LPP systems, as shown in (Figure 20), by using preprocessed video that encodes all 8 viewpoints in each video frame on a per display basis. These viewpoints are then reconstructed on the fly via GPU based interleaving. Playback timing and synchronization is controlled by CGLX [58] and developed on top of the VideoBlaster framework [59]. This approach enables arbitrary-sized tiled LPP video playback.

REVE is a step toward an important goal of 3D visualization technology: making it as easy and convenient to view computational data in 3D as it is to view the data in 2D on TV. However, creating a fully-surround VR system from LPP displays needs significant further research and development.

4.2. Narrow-Bezel Panel-based Systems

As mentioned above, zero-width bezels would make tiled displays more attractive. This section describes some results with narrow-bezel displays designed for multi-tile signage in public spaces.




Figure 24. The Alioscopy displays used in the REVE are meant to be viewed more orthogonally than in the above photographs. A cube-shaped room-sized CAVE made from LCD panels instead of projectors presents problems, such as off-axis intensity attenuation and ghosting, as well as, in this case, depth limitation. Content: UCSD Protein Data Bank. Photo: Tom DeFanti. (Stereo not turned off.)



Figure 25. EVL's CyberCommons and SAGE. Photo: Lance Long.

4.2.1. AESOP (Almost Entirely Seamless OptiPortal)

LCD panels typically have ~18–100 mm borders (called *bezels*) that, when built into tiled displays, create a window mullion-like effect, covering portions of displayed text, data details, or other key features of large, high-resolution images (Figures 9, 10). Eliminating the bezels has long been an obvious need. Design objectives for state-of-the-art tiled display walls are to:

- Create an almost entirely seamless display out of flat panels
- Provide high-resolution and contrast
- Show mono or anaglyph stereo images while awaiting availability of narrow-bezel active or passive stereo displays
- Scalability to large-form-factor collaborative workspaces

Recent commercial offerings by NEC [60] and Samsung [61] address the digital signage market with 1366x768 WXGA ~1-Megapixel “ultra-narrow bezel” displays; 1080p ~2-Megapixel displays are possible by late 2010. Using 46” NEC X461UN LCD 720p (~3.65 mm-bezel) displays, several tiled display configurations were developed. EVL’s 18-tile (3x6) version, produced first, is called *Cyber-Commons* (Figure 25), Calit2’s 4x4 and KAUST’s 4x10 versions are called *AESOP* (Almost Entirely Seamless OptiPortal) (Figure 26). These displays feature inter-tile borders that are ~7 mm thick (the approximate thickness of a standard pencil) when tiled edge-to-edge within the framing, almost eliminating the “window pane” effect of previous OptiPortals, albeit at lower resolution per panel. The 7 mm bezel provides a display quite adequate for PowerPoint presentations or teleconferencing, whereas ~35 mm thick-bezel OptiPortals unacceptably block letters in text slides, scientific information in charts, and important features in faces (Figure 10).

Although the AESOP and Cybercommons displays do not support stereo visuals (other than, of course, anaglyph), the AESOP display has been fitted with a tracker to test the utility and desirability of large-screen imagery presented at the proper perspective for the tracked viewer. (This is, of course, what anyone without functioning stereo vision sees in a CAVE; monovision people get a VR experience in CAVEs from the other 3D cues presented, like proper perspective). AESOP and the NexCAVE provide means to test and compare the value of stereo tracked visuals to mono tracked visuals in various tasks.




Figure 26. KAUST's AESOP showing content by Helmut Pottman. Photo: Tom DeFanti.




Figure 27. The proposed NG-CAVE will be built from tiling 51 near-seamless LCD panels, each fitted with micropolarization overlays providing a continuous circularly polarized 3D screen that can also be used to display 2D information. The NG-CAVE as currently conceived will provide 53 Megapixels of stereoscopic resolution and 106 Megapixels of monoscopic resolution. Illustration: Jason Leigh.

4.2.2. Next-Generation CAVE (NG-CAVE)

Clearly, a micropolarized ultra-narrow bezel display implementation can help minimize the NexCAVE's bezel issues. This is the goal of EVL's proposed NG-CAVE, a larger, more seamless NexCAVE which will be built from ultra-narrow bezel 1080p displays, not yet available but anticipated (Figure 27).

The NG-CAVE will be arranged as a semi-cylindrical configuration that provides a virtually seamless panorama with a horizontal field of view of approximately 180°. A total of 51 display monitors constitute the NG-CAVE, created by tiling 15 LCD panels per side, with 2 columns of LCDs at the corners of the walls. The extra screens provide more surround imagery. The use of micropolarized stereo does not degrade the 2D image when stereo is turned off (text is sharp, lines are not interlaced), allowing the NG-CAVE, to be a dual-purpose display for both 2D and 3D imagery. EVL plans to investigate custom micropolarization strategies to correct for above/below

screen viewing angles, as well as possibly the use of seamless active stereo displays.

5. Software

Classic CAVEs have, from the beginning, used a driver package called CAVELib, originally written by graduate students at EVL, and now commercialized by Mechdyne [62]. Applications in the StarCAVE, NexCAVE and AESOP can use OpenCover, which is the OpenSceneGraph-based VR renderer of COVISE [63]. The NexCAVE and AESOP displays also use CGLX. The Cyber-Commons display runs SAGE; the software for the NG-CAVE has not been determined yet. KAUST also runs AVIZO [64] on its VR displays.

5.1. CGLX

The Cross-Platform Graphics Library (CGLX) provides a common visualization platform, supporting networked, scalable, multi-tile 2D and 3D visualization environments, combined with built-in streaming capabilities. CGLX is a flexible and transparent OpenGL-based graphics framework for distributed, high-performance visualization systems. The framework allows OpenGL based applications to utilize massively scalable visualization clusters, such as multi-projector or high-resolution tiled display environments, and to maximize the achievable performance and resolution.

As such, CGLX combines network-centric scalability with native performance, hardware-accelerated rendering, while exposing an open programming interface (API). It also provides a hardware abstraction and device interface layer, allowing users to integrate interaction devices and paradigms for use in these collaborative digital workspaces. A good example are personal electronics devices that have become a natural part of daily life, including smart phones, media players, and interfaces such as game controllers, and many similar devices. One interface particularly appealing is smart-phone multi-touch (Figure 28).

5.2. SAGE

The Scalable Adaptive Graphics Environment (SAGE) [65] (Figure 24) is a graphics streaming architecture that enables users to interactively access, display and share a variety of data-intensive information, in a variety of resolutions and formats, from multiple sources, with the same ease that the Web affords for accessing lower-resolution objects today. SAGE is cross-platform,




Figure 28. NexCAVE running CGLX, controlled by multi-touch iPhone application, showing LIDAR-scanned image constructed by Vid Petrovic. Photo: Tom DeFanti.

open-source middleware that enables users worldwide to have a common operating environment, or framework, to access, stream and juxtapose data objects, whether digital cinema animations, high-resolution images, high-definition video-teleconferencing, presentation slides, documents, spreadsheets or laptop screens, on one or more tiled display walls. SAGE and tiled display walls are creating a global collaborative visualization environment that allows virtual teams of researchers to manage the scale and complexity of their data and work with one another.

SAGE Visualcasting supports global collaboration by enabling two or more users to share application content, sending multi-gigabit streams as required. Connected, participating endpoint sites form a virtual laboratory, as Visualcasting enables everyone to see the same content at the same time. Endpoints can be of any size and configuration, varying from a single high-resolution monitor to room-sized tiled display walls. Each site maintains control of the layout (position, size, overlapping) of its displays. Visualcasting lets users select what information they want to send, and to whom. Unlike multicast, which requires network engineering, Visualcasting is application-centric.

SAGE supports several mid-air user interface devices; notably, the Gyromouse, joysticks, trackballs, 6 degree-of-freedom magnetic trackers and the Nintendo Wiimote. Multiple devices (not necessarily of the same kind) can simultaneously interact with any of the applications on the display to start/stop applications, manipulate windows (move, resize, maximize, minimize) and interact with applications or their user interfaces (Figure 29).

A full-screen touch interface for the EVL CyberCommons has been added (Figure 30).




Figure 29. SAGE Gyromouse controller at EVL. Photo: Lance Long.



Figure 30. The EVL Cyber-Commons is shown fitted with a multi-user touch screen interface that enables multiple users in front of the wall to interact with the content simultaneously. Photo: Lance Long.

5.3. COVISE

COVISE (Collaborative Visualization and Simulation Environment) is a scientific visualization framework which integrates simulations, post-processing and visualization in one application. COVISE is designed to allow for collaboration between multiple sites, and to integrate CPU and visualization resources at different sites into one system in a transparent way, which makes it as easy to work with such remote resources as it is to work with local ones. COVISE is based on a visual programming paradigm, which allows the user to concatenate simulation, computation, and visualization modules with a graphical user interface, thus creating a data flow network. These modules are implemented as separate processes, which can run on the local machine or on any other COVISE system on the network. Several video segments are on the Calit2 website [66].

One of COVISE's visualization modules is OpenCOVER, a sophisticated virtual reality rendering module, which can run within a COVISE module

pipeline, or stand-alone. OpenCOVER supports virtual environments like head-tracked single-screen systems, PC cluster-based systems, powerwalls, curved screens, dome systems and CAVEs. OpenCOVER also supports most of the interaction devices commonly used. Once the configuration of the visualization hardware has been set up in COVISE, the user can switch between different visualization hardware without the need to adjust the visualization module. OpenCOVER provides full VRML97 capabilities, which allows application designers to use commercial 3D modeling tools like 3ds Max to create interactive virtual environments, and it supports sound effects. OpenCOVER also provides a C++ and OpenSceneGraph-based plug-in system for the application developer, which exposes the rendering and interaction interface to the programmer, allowing for full control of the virtual environment, without the need to modify COVISE-internal source code. While the authors use COVISE exclusively on Linux systems, it is also compatible with Windows and MacOS.

6. Conclusions

CAVEs, derivatives, and similar screen-based, large-scale VR devices, exploiting technology advances and commoditization, have been providing users steadily higher resolution, better contrast, and improvements in 3D stereo graphics, texture and image mapping, and functional networked tele-immersive collaboration. As CAVEs of the future increasingly adopt panel-based technology, normal-ceilinged lab, office, work, greeting and living spaces can house walk-in VR displays. At the same time, the recent market availability of consumer 3D HDTVs has dropped the retail costs of the displays by two-thirds or more. PC motherboard/GPU card configurations are becoming available that can drive a dozen or more panels, so the CAVE, perhaps by its 21st birthday, is destined to become a “prosumer” visualization device. Achieving all the conceivable goals for a perfect CAVE will inspire and, indeed, require at least another decade of research; the future of the CAVE is bright.

Acknowledgements

This publication is based on work supported in part by Award No US-2008-107/SA-C0064, made by King Abdullah University of Science and Technology (KAUST). Mr. Sami Maghlooth was the project manager for the KAUST visualization facilities design and construction phase.

UCSD, through Calit2, received major funding from the State of California for the StarCAVE. UCSD receives major funding from the National Science Foundation (NSF), awards CNS-0821155, which supported the construction of the Calit2 AESOP display.

UIC receives major funding from the National Science Foundation (NSF), awards CNS-0420477 (LambdaVision), OCI-0943559 (SAGE) and CNS-0959053 (NG-CAVE), and also receives support from Sharp Laboratories of America and the State of Illinois. Also, UCSD and UIC were lead institutions of the NSF-funded OptIPuter project, award OCI-0225642.

EVL VR work during the past 20 years has been supported by numerous NSF, DARPA, and DOE awards to UIC. The UIC/EVL versions of CAVELib software were authored by Carolina Cruz-Neira and Dave Pape. Maggie Rawlings and Jim Angelilo were key to the commercialization of the CAVE; the royalties supported many students and staff members at EVL.

Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the funding agencies and companies.

Avizo® is a registered trademarks of VSG, Visualization Sciences Group SAS. CAVE™ is a trademark of the Board of Trustees of the University of Illinois.

References

- [1] Cruz-Neira C., Sandin D., DeFanti T., *et al.*, The CAVE, Communications of the ACM, 35(6), 1992, 64-72
- [2] Kenyon R. V., Sandin D. Smith, Randall C., Pawlicki R., *et al.*, Size-Constancy in the CAVE, Presence: Teleoperators and Virtual Environments, 16(2), 2007, 172-187
- [3] <http://en.wikipedia.org/wiki/Cave%5Automatic%5Virtual%5Environment>
- [4] Sutherland I. E., A head-mounted three dimensional display, In: Proceedings of Fall joint computer conference, part I, ACM (December 9-11, 1968), 1968, 757-764
- [5] Tan D., Gergle D., Scupelli P. and Pausch R., With similar visual angles, larger displays improve spatial performance, In: Proceedings of the SIGCHI conference on Human factors in computing systems, 2003
- [6] Tan D., Gergle D., Scupelli P. and Pausch R., Physically large displays improve performance

- on spatial tasks, In: ACM Transactions on Computer-Human Interaction, 2006
- [7] Yost B., Hacıahmetoglu Y. and North C., Beyond visual acuity: the perceptual scalability of information visualizations for large displays, In: Proceedings of the SIGCHI conference on Human factors in computing systems, 2007
 - [8] Korab H., Brown M., (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, <http://www.ncsa.uiuc.edu/General/Training/SC95/GII.HPCC.html>
 - [9] DeFanti T., Brown M., and Stevens R. (Guest Editors), Virtual Reality Over High-Speed Networks, IEEE Computer Graphics & Applications, 16(4), 1996, 14-17, 42-84
 - [10] Lehner V.D., DeFanti T., Distributed Virtual Reality: Supporting Remote Collaboration in Vehicle Design, IEEE Computer Graphics & Applications, 1997, 13-17
 - [11] Leigh J., Johnson A., DeFanti T., CAVERN: A Distributed Architecture for Supporting Scalable Persistence and Interoperability in Collaborative Virtual Environments, Virtual Reality: Research, Development and Applications, 1997
 - [12] Leigh J., DeFanti T., Johnson J., Brown M., *et al.*, Global Tele-Immersion: Better Than Being There, In: ICAT '97, 7th Annual International Conference on Artificial Reality and Tele-Existence (December 3-5, 1997, Virtual Reality Society of Japan, University of Tokyo, Japan), 1997, 10-17
 - [13] Leigh J., Park K., Kenyon R.V., Johnson A.E., *et al.*, Preliminary STARTAP Tele-Immersion Experiments between Chicago and Singapore, In: 3rd High Performance Computing Asia Conference & Exhibition (September, 1998, Singapore), 22-25, 1998, 687-693
 - [14] Leigh J., Johnson A., Renambot L., DeFanti T., *et al.*, Emerging from the CAVE: Collaboration in Ultra High Resolution Environments, In: Proceedings of the First International Symposium on Universal Communication (June, 14-15 2007, Kyoto, Japan)
 - [15] Johnson A., Leigh J., Costigan J., Multiway Tele-Immersion at Supercomputing '97, IEEE Computer Graphics and Applications, 1998
 - [16] Stevens R., DeFanti T. Tele-Immersion and Collaborative Virtual Environments, In: The Grid: Blueprint for a New Computing Infrastructure, I. Foster and C. Kesselman (Eds.), Morgan Kaufmann Publishers, 1999, 131-158
 - [17] Park K., Cho Y., Krishnaprasad N., Scharver C., *et al.*, CAVERNsoft G2: A Toolkit for High Performance Tele-Immersive Collaboration, In: Proceedings of the ACM Symposium on Virtual Reality Software and Technology 2000 (October, 2000, Seoul, Korea)
 - [18] Smarr L., The Emergence of a Planetary-Scale Collaboratory for Data-Intensive Research, Points of View - a tribute to Alan Kay, Ian Piumarta and Kim Rose (Eds.), Viewpoints Research Institute, 2010, 79-96
 - [19] DeFanti T., Sandin D., Brown M., Pape, *et al.*, Technologies for Virtual Reality/Tele-Immersion Applications: Issues of Research in Image Display and Global Networking (chapter), Frontiers of Human-Centered Computing, Online Communities and Virtual Environments, Rae Earnshaw, Richard Guedj, Andries van Dam and John Vince (editors), Springer-Verlag London, 2001, 137-159
 - [20] DeFanti T., Leigh J., Brown M., Sandin, *et al.*, Teleimmersion and Visualization with the OptIPuter, In: Telecommunication, Teleimmersion and Telexistence, (Susumu Tachi, editor), Ohmsha/IOS Press, 2003, 25-71
 - [21] Leigh J., Johnson A., Supporting Transcontinental Collaborative Work in Persistent Virtual Environments, IEEE Computer Graphics and Applications, 1996
 - [22] Leigh J., Johnson A., Vasilakis C., DeFanti T., Multi-Perspective Collaborative Design in Persistent Networked Virtual Environments, In: Proceedings of the IEEE Virtual Reality Annual International Symposium VRAIS '96, (March 1996, Santa Clara, CA)
 - [23] Leigh J., Johnson A., DeFanti T., Issues in the Design of a Flexible Distributed Architecture for Supporting Persistence and Interoperability in Collaborative Virtual Environments, In: Proceedings of Supercomputing '97, (November, 15-21, 1997, San Jose, CA)
 - [24] Brown M., DeFanti T., *et al.*, The International Grid (iGrid): Empowering Global Research Community Networking Using High Performance International Internet Services, Proceedings of INET '99, (June, 22-25, 1999, San Jose, CA)
 - [25] Jeong B., Renambot L., Jagodic R., Singh R., *et al.*, High-Performance Dynamic Graphics Streaming for Scalable Adaptive Graphics Environment, ACM/IEEE Supercomputing 2006, (November 11-17, 2006)
 - [26] Smarr L., Brown M., DeFanti T., de Laat C. (guest editors), Special issue on the International Grid (iGrid) 2005 Workshop, Future Generation Computer Systems/The International Journal of Grid Computing: Theory, Methods and Applications, Elsevier B.V., 22(8), (October, 01-31, 2006)

- [27] <http://www.glif.is>
- [28] Harrison C., Tan D., Morris D., Skinput: Appropriating the Body as an Input Surface, CHI 2010, (April 10–15, 2010, Atlanta, Georgia, USA)
- [29] DeFanti T., Dawe G., Sandin D., Schulze J., *et al.*, The StarCave, a third-generation cave and virtual reality Optiportal, In: Future Generation Computer Systems/The International Journal of Grid Computing: Theory, Methods and Applications, Elsevier B.V., 25(2), 2008, <http://ivl.calit2.net/wiki/index.php/StarCAVE%5Measurement%5Diagrams>
- [30] SXRD 4K Ultra-high Resolution Projectors, Sony product manual, 2009, <http://pro.sony.com/bbsc/ssr/micro-sxrdsite/>
- [31] <http://www.meyersound.com/pdf/case%5studies/cs%5kaust.pdf>
- [32] Rhee J., Schulze J.P., DeFanti T., Hotspot Mitigation in the StarCAVE, In: Proceedings of SPIE Vol. 7525: The Engineering Reality of Virtual Reality, San Jose, CA, 2010
- [33] <http://www.samsung.com/au/tv/warning.html>
- [34] Lipton L., The CrystalEyes Handbook. San Rafael, CA: StereoGraphics Corporation, 1991
- [35] Sandin D., Sandor E., Cunnally W., Resch, M. *et al.*, Three-Dimensional Visualization and Display Technologies, Proc. of SPIE Vol. 1083, 1989
- [36] Sandin D., Margolis T., Dawe G., Leigh J., *et al.*, The Varrier™ Auto-Stereographic Display, In: Proceedings of Photonics West 2001: Electronics Imaging, sponsored by SPIE, (January 20–26, 2001, San Jose, CA), <http://spie.org/web/meetings/programs/pw01/home.html>
- [37] Sandin D., Margolis T., Ge J., Girado J., *et al.*, The Varrier Autostereoscopic Virtual Reality Display, ACM Transactions on Graphics, In: Proceedings of ACM SIGGRAPH 2005, (July 31– August 4, 2005, Los Angeles, CA), 24(3), 2005, 894–903
- [38] Peterka T., Kooima R.L., Girado J.I., Ge J., *et al.*, Evolution of the Varrier autostereoscopic VR display, In: Proceedings IS&T / SPIE Electronic Imaging 2007, San Jose, CA, 2007
- [39] Peterka T., Kooima R.L., Girado J.I., Ge J., *et al.*, Dynallax: Solid State Dynamic Parallax Barrier Autostereoscopic VR Display, In: Proceedings of the IEEE Virtual Reality Conference 2007 - VR'07, (March, 10–14, 2007, Charlotte, NC)
- [40] Peterka T., Kooima R.L., Sandin D., Johnson A., *et al.*, Advances in the dynallax solid-state dynamic parallax barrier autostereoscopic visualization display system, IEEE Transactions on Visualization and Computer Graphics, 14(3), 2008, 487–499
- [41] Peterka T., Ross R., Yu H., Ma K.-L., Autostereoscopic display of large-scale scientific visualization, In: Proceedings IS&T / SPIE Electronic Imaging 2009, San Jose, CA, 2009
- [42] <http://en.wikipedia.org/wiki/20/20%5vision#Visual%5acuity%5expression>
- [43] DeFanti T., Leigh J., Renambot L., Jeong B., *et al.*, The OptiPortal, a Scalable Visualization, Storage, and Computing Interface Device for the OptiPuter, In: Future Generation Computer Systems/The International Journal of Grid Computing: Theory, Methods and Applications, Elsevier B.V., 25(2), 2008
- [44] <http://www.optiputer.net>
- [45] Johnson A., Sandin D., Dawe G., Qiu Z., *et al.*, Developing the PARIS: Using the CAVE to Prototype a New VR Display, In: Proceedings of IPT 2000, (June 19–20, 2000, Ames, IA), <http://www.evl.uic.edu/EVL/RESEARCH/paris.shtml>
- [46] <http://www.calit2.net/newsroom/release.php?id=1708>
- [47] Faris S.M., Method and System for Producing Micropolarization Panels for Use in Micropolarizing Spatially Multiplexed Images of 3-D Objects During Stereoscopic Display Processes. United States Patent 5844717, 1998, <http://www.patentstorm.us/patents/5844717/claims.html>
- [48] Benton S.A., Slowe T.E., Kropp A.B., and Smith S.L., Micropolarizer-based multiple-viewer autostereoscopic display, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Presented at the Society of Photo-Optical Instrumentation Engineers (SPIE) Conference, Vol. 3639, 1999, 76–83
- [49] Arisawa Mfg. Co., Ltd., 2009, <http://www.arisawa.co.jp/en/product/3d.html>
- [50] Kooima R., Leigh J., Johnson A., Roberts D., *et al.*, Planetary-Scale Terrain Composition, IEEE Transactions on Visualization and Computer Graphics, IEEE Computer Society, April, 2009
- [51] Sandin D., Margolis T., Dawe G., Leigh J., *et al.*, The Varrier Auto-Stereographic Display, The Stereoscopic Displays and Virtual Reality Systems IIX, (January 24, 2002, San Jose, CA)
- [52] Girado J., Real-Time 3D Head Position Tracker System with Stereo Cameras Using a Face Recognition Neural Network, PhD thesis, University of Illinois at Chicago, 2004
- [53] Fukushima R., Taira K., Saishu T., Momonoi Y., *et al.*, Effect of light ray overlap between neighboring parallax images in autostereoscopic 3D displays, In: Proceedings SPIE, Vol. 7237
- [54] Zwicker M., Matusik W., Durand F., Pfister H., Antialiasing for Automultiscopic 3D Displays,

- Eurographics Symposium on Rendering, 2006
- [55] Kooima R., Prudhomme A., Schulze J., Sandin D., *et al.*, A Multi-viewer Tiled Mutostereoscopic Virtual Reality Display, To appear in the Proceedings of the 17th ACM Symposium on Virtual Reality Software and Technology, (November, 2010, Hong Kong)
 - [56] <http://www.alioscopy.com/>
 - [57] Girado J., Sandin D., DeFanti T., Wolf L., Real-Time Camera Based Face Detection using a modified LAMSTAR Neural Network System, In: Proceedings of Electronic Imaging 2003/Conference 5015: Applications of Artificial Neural Networks in Image Processing VIII, 15th Annual Symposium on Electronic Imaging: Science and Technology, sponsored by The Society for Imaging Science and Technology (IS&T) and The International Society for Optical Engineering (SPIE), (January 20-24, 2003, Santa Clara, California, USA)
 - [58] Doerr K.U., Kuester F., CGLX: A Scalable, High-performance Visualization Framework for Networked Display Environments, IEEE Transactions on Visualization and Computer Graphics, Vol. 99, 2010
 - [59] Ponto K., Wypych T., Doerr K., Yamaoka S., *et al.*, VideoBlaster: A Distributed, Low-Network Bandwidth Method for Multimedia Playback on Tiled Display Systems, 11th IEEE International Symposium on Multimedia, 2009, 201-206
 - [60] <http://www.necddisplay.com/NewTechnologies/UltraNarrowLCD/>
 - [61] <http://www.samsung.com/us/consumer/professional-displays/professional-displays/lcd/LH46MVTB/N/ZA/index.idx?pagetype=prd%5detail>
 - [62] <http://www.mechdyne.com/integratedSolutions/software/products/CAVELib/CAVELib.htm> and <http://www.vrco.com/CAVELib/OverviewCAVELib.html>
 - [63] Rantzau D., Lang U., Ruehle R., Collaborative and Interactive Visualization in a Distributed High Performance Software Environment, In: Proceedings of the International Workshop on High Performance Computing for Graphics and Visualization, Swansea, Wales, 1996, <http://www.hlrs.de/organization/vis/covise/features/opencover/>
 - [64] <http://www.vsg3d.com/vsg%5prod%5avizo%5overview.php>
 - [65] <http://www.sagecommons.org>
 - [66] <http://www.calit2.net/newsroom/article.php?id=1584>
<http://www.youtube.com/calit2#p/search/0/pFTfZ0VtzIU> (NexCAVE display of 3D model of rebar layout for new span of the San Francisco Bay Bridge
<http://www.youtube.com/calit2#p/search/1/9tEkHz97yNU> (NexCAVE display of 3D model of Calit2 headquarters building at UC San Diego)
<http://www.youtube.com/calit2#p/search/2/AEjJYepZ%5FP8> (NexCAVE display of archeological dig site model)
<http://www.youtube.com/calit2#p/search/3/Ze5IprHtC5c> (NexCAVE with 3D model of de-salination plant designed at the National University of Singapore)
<http://www.youtube.com/calit2#p/search/4/cQP%5YhtAeyQ> (NexCAVE exploration of 3D model for the holy shrine at Mecca)
<http://www.youtube.com/calit2#p/search/5/u96M4OhwIgA> (NexCAVE Exploration of Jordan Archaeological Excavation Site)