

Janne Jalkanen:

Building A Spatially Immersive Display: HUTCAVE



Lisensiaattityö / Licentiate Thesis

Työn valvoja: Prof. Tapio Takala
Työn ohjaaja: Prof. Tapio Takala

Espoossa 4 maaliskuuta 2000.

Author:	Janne Jalkanen	
Title of the Thesis:	Building a spatially immersive display - HUTCAVE	
Title in Finnish:	Keinotodellisuuslaitteiston rakentaminen	
Date:	4 March 2000	Number of Pages: 132
Department:	Department of Computer Science	Chair: Telecommunications Software and Multimedia.
Supervisor:	Prof. Tapio Takala	
Instructor:	Prof. Tapio Takala	
<p>A spatially immersive display is a display that surrounds the user, thus removing or alleviating many disadvantages the common virtual reality systems, such as head-mounted displays have. The most common example of these spatially immersive displays is the CAVE, "CAVE Automatic Virtual Environment", first built at University of Illinois, in 1993. It combines a large field-of-view with high-resolution images and a high frame refresh rate.</p> <p>In this work, the current Virtual Reality (VR) and Virtual Environment (VE) systems are examined, and then the CAVE construction is presented. Principles of stereo vision are explained and current methods of obtaining both autostereoscopic and stereopsis-based vision are reviewed.</p> <p>Aspects of different projection methods, screens, mirrors, projectors, tracking equipment, and computing systems are examined. Also, recent work in CAVE audio, so far neglected in research, is presented. Some of the mathematics is also explained, since in most CAVE-systems some sort of optical folding is necessary.</p> <p>Two cases of CAVE construction are presented, both at the Helsinki University of Technology. The first is a single-wall installation built as a temporary system, and the second is a four-sided CAVE at a new location, superseding the temporary installation.</p> <p>Finally the conclusions are presented, both from the process management point of view, and from the technical point of view, examining the good and bad points of the chosen solutions.</p>		
Keywords:	CAVE, virtual reality, virtual environments	
Not borrowable till:	Library code:	

Tekijä:	Janne Jalkanen		
Työn nimi:	Keinotodellisuuslaitteiston rakentaminen		
Title in English:	Building a spatially immersive display - HUTCAVE		
Päivämäärä:	4 maaliskuuta 2000	Sivumäärä:	132
Osasto/Laitos:	Tietotekniikan osasto	Professuuri:	Tietoliikenneohjelmistot ja multimedia
Työn valvoja:	Prof. Tapio Takala		
Työn ohjaaja(t):	Prof. Tapio Takala		
<p>Spatiaalisesti immersiiivinen näyttö (Spatially Immersive Display, SID) on keinotodellisuusnäyttö, joka pyrkii ympäröimään katsojan, ja näin poistaa tai vähentää useimpia ongelmia, joita esim. kannettaviin keinotodellisuusnäyttöihin liittyy. Yleisimmin tunnettu esimerkki immersiiivisistä näytöistä on virtuaalihuone, eli CAVE, jollainen rakennettiin ensimmäisen kerran vuonna 1993 Illinoisin yliopistossa. Siinä yhdistyvät laaja näkökenttä, korkearesoluutioiset kuvat ja nopea virkistystaajuus.</p> <p>Tässä työssä tutustutaan ensin nykyisiin keinotodellisuus- ja keinoympäristöjärjestelmiin ja esitellään CAVE. Stereonäköön liittyvät periaatteet ja nykyiset menetelmät syvyyseffektin aikaansaamiseksi esitellään.</p> <p>Tämän jälkeen tutustutaan eri projektiomenetelmiin, projisiopintoihin, projekto-reihin, peileihin, paikannuslaitteisiin ja tietokonejärjestelmiin. Äänipuolta käydään myös läpi, sillä monessa CAVEssa tähän ei ole kiinnitetty suurempaa huomiota. Myös taitettuun optiikkaan liittyvä matematiikka käydään läpi.</p> <p>Kaksi esimerkkitapausta CAVEN rakentamisesta käydään läpi. Molemmat CAVEt rakennettiin TKK:lle, ensin väliaikaisena järjestelmänä yksiseinäisenä, ja sitten neliseinäisenä versiona uudessa rakennuksessa.</p> <p>Lopuksi käydään läpi johtopäätöksiä sekä tekniseltä, että myös projektihallinnolliselta näkökannalta.</p>			
Avainsanat:	Virtuaalihuone, keinotodellisuus, keinoympäristöt		
Ei lainata ennen:	Työn sijaintipaikka:		

This page intentionally left blank.

PREFACE¹

"Who are you? What do you want? Why are you here? Where are you going?"
"Lorien."
"Did you think we had forgotten you? We have been waiting .. for you."
"Beyond the rim?"
"Yes."
"There is .. so much I still don't understand."
"As it should be."
"Can I come back?"
"No. This journey has ended. Another begins. Time .. to rest now."
— *Lorien and Sheridan in Babylon 5: "Sleeping in Light"*

This has truly been a learning experience.

Thanks go to (in no particular order) to people who contributed to this thesis:

- Erik Bunn for projectorUI, the MatLab script responsible for fitting everything in.
- Matti Gröhn for support and, of course, his bottomless Single Malt supply.
- Helsinki Graduate School in Computer Sciences (HeCSE) for supporting the work in this thesis by granting me a scholarship.
- Professor Tassu "Tapio" Takala, for letting me work on such an interesting project.
- Hannu Napari for his help on the hardware - especially all those hours spent adjusting the projector - and the proofreading of this thesis.
- DIVA Group (Lauri, Tapio, Jarmo, Tommi) for acoustics work and help.

And finally, thanks to Heli, the love of my life.

Janne Jalkanen
Espoo, Finland, 4 March 2000
Janne.Jalkanen@iki.fi

Faith manages.

1. The CAVETM is a trademark of the Board of Trustees of the University of Illinois (UIC). All other products mentioned are registered trademarks or trademarks of their respective holders. No harm is intended by using them in this scientific paper.

This page intentionally left blank.

LIST OF USED ABBREVIATIONS

2D	Two-Dimensional
3D	Three-Dimensional
AR	Augmented Reality
CAVE	CAVE Automatic Virtual Environment
CPU	Central Processing Unit
CRT.....	Cathode Ray Tube
CS	Computer Science
C ³ I.....	Command, Control, Communications, and Intelligence. Used in military.
DG.....	Display Generator
DLP	Digital Light Processing
DMD.....	Digital Micromirror Device
EVL	Electronic Visualization Laboratory, University of Illinois.
FOV.....	Field Of View
FSS	Field Sequential Stereo, one of different stereo formats.
GE	Geometry Engine™
HMD.....	Head Mounted Display
HRTF	Head-Related Transfer Function
HUT.....	Helsinki University of Technology
IPD	Interpupillary Distance, also known as IOD (inter-ocular distance).
IR	Infrared
LCD.....	Liquid Crystal Display
LEEP.....	Large Expanse, Extra Perspective. The name of one of the first VR displays.
MAVERIK.....	Manchester Virtual Environment Interface Kernel. A freely available VR software.
MR	Mediated Reality / Mixed Reality
NTSC	National Television Systems Committee. TV standard currently in use in USA and Japan.
OS.....	Operating System
PAL	Phase Alternating Lines. TV standard currently in use throughout most of Europe.
PC	Phase Coherent
RM	Raster Manager
SGI.....	Former Silicon Graphics, Inc.
SID	Spatially Immersive Display
SMP	Symmetrical Multi-Processing
TML.....	Telecommunications and Multimedia Laboratory
TOF	Time of Flight
VBAP	Vector Base Amplitude Panning
VE	Virtual Environment
VIEW.....	Virtual Interactive Environment Workstation
VMD	Virtual Model Display
VR	Virtual Reality
WLC.....	White-Line-Code. One of different stereo formats.
ZPS.....	Zero Parallax Setting

This page still intentionally left blank.

CONTENTS

LIST OF FIGURES	ix
1. Introduction.....	1
1.1 Virtual Reality Technologies - Short History	2
1.1.1 Sutherland and the early pioneers	2
1.1.2 LEEP and commercial applications	2
1.1.3 NASA Ames VIEW and the military	2
1.2 Principles of immersivity	3
1.2.1 Immersion and Senses.....	4
1.2.2 The Difficulty of Fooling The Eye	5
1.3 How Three-Dimensional Displays Are Made?.....	8
1.3.1 Autostereoscopic displays.....	8
1.3.2 Stereograms and Head-Mounted Displays	9
1.3.3 Stereographic Displays	10
1.3.4 It Ain't Real, Though.....	11
1.4 Current immersive VR display technology	12
1.4.1 Classification of VR display systems	12
1.4.2 HMD devices	13
1.4.3 Virtual Model Displays.....	14
1.4.4 Spatially Immersive Displays, such as CAVE	16
1.4.5 Futuristic displays.....	20
2. CAVE Automatic Virtual Environment.....	23
2.1 Physical Structure	25
2.1.1 Mirror positioning.....	26
2.1.2 Height & Floor Projection.....	28
2.1.3 Frames.....	28
2.1.4 Fully enclosed CAVEs	30
2.2 Back-projection systems	31
2.2.1 Image Separation.....	31
2.2.2 Screen Type.....	31
2.2.3 Brightness and Contrast	32
2.2.4 Hot Spots and Other Odd Problems	34
2.2.5 Floor	37
2.2.6 Corners and Edges	38
2.2.7 Mirrors	42
2.3 Projectors.....	44
2.3.1 Display Resolution	44
2.3.2 Display Refresh Rate	45
2.3.3 Bandwidth.....	46
2.3.4 Geometrical and Other Distortions	47
2.3.5 Projector Types.....	48
2.4 Tracking Systems	50
2.4.1 Tracking Systems Properties.....	50
2.4.2 Tracking Hardware	51
2.5 Computing Equipment.....	54

2.5.1 Main Graphics Computer	54
2.5.2 Auxiliary Computers	58
2.6 CAVE Software	59
2.6.1 Basic Requirements for Software	59
2.7 Audio Systems	61
2.7.1 Acoustics	61
2.7.2 Sound Reproduction	61
2.7.3 Other Sound Sources	62
2.8 Infrastructure	63
2.8.1 Cabling	63
2.8.2 Layout	64
2.8.3 Miscellaneous	64
3. HUTCave. Phase One	67
3.1 In the beginning	67
3.2 Project history	67
3.3 Computing Equipment	68
3.3.1 Main Computer	68
3.3.2 Auxiliary Computers	69
3.4 Software	69
3.4.1 Software Solutions	69
3.4.2 HUT Software	70
3.5 Choosing the Place	71
3.6 Reconstruction	72
3.6.1 Computer Room	72
3.6.2 Power	72
3.6.3 Acoustics and lighting	73
3.6.4 Ventilation and cooling	74
3.7 The HUTCAVE Projection System	75
3.7.1 The Screen	75
3.7.2 The Frame	75
3.7.3 The Projector	77
3.8 Audio System	77
3.9 Installing HUTCAVE equipment	77
3.10 Miscellaneous	79
3.10.1 Tracking	79
3.10.2 Backups	79
3.10.3 Usage	79
3.11 Comments on Phase I	80
4. HUTCAVE Phase Two	82
4.1 Construction work	82
4.1.1 Appearance	82
4.1.2 Acoustics and Lighting	84
4.1.3 Ventilation and Cooling	86
4.1.4 Power	87
4.1.5 The machine room	88
4.1.6 Other modifications	89
4.2 Moving In	90
4.3 HUTCAVE Projection System	91

4.3.1 Physical constraints	91
4.3.2 Delivery and Configuration	92
4.3.3 Construction	93
4.3.4 Problems with the projection system	97
4.4 The audio system	97
4.5 Miscellaneous	99
4.5.1 Wiring	99
4.5.2 Operators	99
4.6 Operational notes	100
4.6.1 Demos	100
4.6.2 Reservations	100
4.6.3 Screens and Physical Construction	100
4.6.4 HUTCAVE as a working environment	100
5. Conclusions and Future Work	102
5.1 Conclusions	102
5.1.1 Quality of the work	102
5.1.2 Organization and process	103
5.1.3 In retrospect...	103
5.2 Future Work	104
5.2.1 Work on the HUTCAVE	104
5.2.2 Future of Virtual Environments	105
Appendix A: Projector Geometry	107
GLOSSARY	111
REFERENCES	115
INDEX	129

LIST OF FIGURES

Figure 1: The field of vision of the human (from [Kalawsky 93]).	6
Figure 2: Parallax classifications and IPD.	8
Figure 3: Virtual Reality Head Mounted Display, which is freely worn.	14
Figure 4: The BOOM-3C device from FakeSpace, Inc.	14
Figure 5: The ImmersaDesk by Pyramid Systems.	15
Figure 6: The Holobench. Image courtesy of TAN GmbH.	16
Figure 7: A typical 4-sided CAVE structure. Notice the floor projection is from above.	17
Figure 8: The KTH Six-sided CAVE built by TAN GmbH. Image courtesy of PDC/KTH.	18
Figure 9: The PRODAS curved screen from SEOS installed at SGI Reading, UK.	18
Figure 10: The VisionDome from ARC. Image courtesy of Alternate Realities Corporation.	19
Figure 11: The CyberSphere system. Copyright Volker Steiger/Science Photo Library	20
Figure 12: The Holodeck from Star Trek - The Next Generation. Star Trek is a registered trademark of Paramount Pictures.	22
Figure 13: Projector distances.	25
Figure 14: CAVE layout, top view. Projectors on the sides.	26
Figure 15: Single-mirror configuration and parameters.	27
Figure 16: Alternate configurations for the Cave side projectors.	27
Figure 17: Cave layout, side view. Left side shows a typical CAVE with a top projection, the right side shows a CAVE with both floor and ceiling projections.	29
Figure 18: Light spillage from one screen to the next.	33
Figure 19: The absorption of light.	34
Figure 20: Hot spot effect and the projection geometry.	35
Figure 21: The Fresnel lens.	35
Figure 22: EVL/Pyramid Systems CAVE corner solution.	39
Figure 23: Cross Edge solution.	40
Figure 24: Welding or seaming the edges.	41
Figure 25: Joining the (front-projected) floor seamlessly with the walls.	41
Figure 26: Sharp edges	42
Figure 27: CRT projector light cones and convergence.	43
Figure 28: Losing physical pixels due to aspect ratio.	44
Figure 29: Alternate tracker/wiring placements.	66
Figure 30: The Media Brewery facilities where the HUTCave was located in Phase I	71
Figure 31: View from the control room into the CAVE.	73
Figure 32: Left, attaching the screen to the frames. Right, the frame stand.	76
Figure 33: The frame inside the CAVE room, without projector or mirror.	76
Figure 34: HUTCave personnel installing the Onyx2 and the mirror.	78
Figure 35: The HUTCAVE in Phase I configuration running an architectural application.	79
Figure 36: The DIVA Virtual Orchestra running in the Phase I HUTCAVE.	80
Figure 37: The HUTCave phase I running.	81
Figure 38: The CAVE room layout.	83
Figure 39: The CAVE room layout, side view.	84
Figure 40: Installing the acoustic panels inside the Cave room.	85
Figure 41: Covered windows (left) and the door (right).	86
Figure 42: The extra pump for condensation water.	89
Figure 43: The Onyx2 installed. On the right side of the Onyx, an ASCII console for controlling the computer in case of problems, and the backup devices.	91
Figure 44: The Matlab projectorUI script User Interface. Image courtesy of Erik Bunn.	92

Figure 45: The floor and the top frame.	93
Figure 46: The front of the frame structure in place.	94
Figure 47: Detail of the front corner. Notice the bar around which the front screen is wrapped. 95	
Figure 48: The top left corner. Notice how the vertical and horizontal support bars are attached to the screen itself.	95
Figure 49: Detail of the front left corner from the outside.	96
Figure 50: The mirror from the side.	96
Figure 51: Loudspeakers in the ceiling. Note the two slightly differing ways of mounting them. 98	
Figure 52: The floor and head-level mountings of the loudspeakers.	98
Figure 53: Two operators running the CAVE tests.	99
Figure 54: Projector geometry, viewed from top.	107
Figure 55: Image geometry on the mirror, viewed from the side.	109

1. Introduction

What is *Virtual Reality*? According to Kalawsky [Kalawsky 93], in every conference someone attempts to define their version of “Virtual Reality”, and then comments that “there probably are as many definitions for virtual reality as there are people in the field.” When one adds the confusion that exists in the mainstream media about the whole term, and one has something that everyone sort of knows of, but nobody can claim to truly, really *know* [Bierbaum 98].

For the purpose of this thesis, I will use the term Virtual Reality for what is most commonly termed as *Virtual Environment*, that is, a synthetically created environment, which allows interaction with it. The perception of this environment can be done via visible, audible, or tactile means, in all cases using the full extent of the senses: stereographic vision for the eyes, surround sound for the ears.

For clarity, the often used term *Augmented Reality* [Caudell 92] should also be explained: Augmented Reality (AR) adds new information to the normal world which we perceive around us. The broad definition as given by Milgram goes: “AR covers any case in which an otherwise real environment is ‘augmented’ by means of virtual (computer graphic) objects” [Milgram 99]. One should notice that this encompasses all large screen, monitor-based, and personal displays.

Augmented Reality is sometimes considered to be a subset of *Augmented Virtuality* (AV) [Milgram 99], where computer-generated scenes are augmented using real-world images, sounds, movies, etc. This is much closer in both spirit and technology to Virtual Environments, even though the applications may slightly differ. Sometimes the term *Mixed Reality* is used to mean the synthesis of AR and AV [Sato 99], [Tamura 99].

Steve Mann has introduced another concept, *Mediated Reality* [Mann 94], which is based on the concept of filtering the reality that is available to us. It includes Augmented Reality, but it also is concerned with removing information (i.e. Diminished Reality) or changing the information around us. For example, a pair of eyeglasses could remove indecent material, displaying pictures of flower bushes instead [Mann 98].

Structure of This Licentiate Thesis

This thesis has been structured into three major parts: First, I will examine the history and the general technologies behind Virtual Reality and related fields in Chapter 1. Second, I will take a look at a single, popular VR installation, the CAVE, and especially its technology in Chapter 2., and finally, I will discuss the actual construction of the HUTCAVE at the Helsinki University of Technology in Chapters 3. and 4., with a set of conclusions to follow in Chapter 5.

The three parts may be read independently of each other, but to fully understand the technology, they should be read in order.

1.1 Virtual Reality Technologies - Short History

1.1.1 Sutherland and the early pioneers

According to [Kalawsky 93], the first system that could be considered to be VR was the Philco Corporation Headsight television surveillance system, which allowed an observer to view dangerous operations from a safe, remote location. This system was already produced in 1961, almost seven years before the Ivan Sutherland's Ultimate Display which is often regarded as the starting point of Virtual Reality hardware. Of course, the Headsight can be considered more as telepresence instead of Virtual Reality, and thus more of an Augmented Reality-type application.

The first undisputed VR installation was Dr. Ivan Sutherland's Ultimate Display, which consisted of two miniature Cathode Ray Tubes (CRT) mounted on a head band, which was then attached to a mechanical arm which could sense the position and orientation of the user's head, aided by an ultrasonic tracker. The system allowed the user to see wire-frame 3D objects that could be overlaid on the real world, in a form of Augmented Reality.

The Grope system from University of North Carolina (1967) was yet another step forward, as it introduced the haptic display, that is, a system which allowed the user to actually feel things inside the virtual space. According to the early results, the haptic feedback proved to improve the feeling of presence, even better than stereo vision.

1.1.2 LEEP and commercial applications

The Large Expanse, Extra Perspective (LEEP) optical system was designed by Eric Howlett in 1975 and provides the basis for most of the current virtual reality helmets available today. The original LEEP system was redesigned for the NASA Ames Research Center in 1985 for their first virtual reality installation.

The DataGlove was born at about the same time, the first glove-like device being designed at the University of Illinois, Chicago in 1977 by Dan Sandin, Richard Sayre and Thomas DeFanti. However, it wasn't until the establishment of the VPL Research, their DataGlove, and the article in Scientific American [Foley 87] that the world became very excited by "Virtual Reality".

1.1.3 NASA Ames VIEW and the military

The VIEW (Virtual Interactive Environment Workstation) was designed in 1985 in the NASA Ames Research Center by Scott Fisher. The system was built according to lessons learned using the LEEP display earlier, and proved to be quite impressive. It already featured many techniques that are used nowadays: a Polhemus tracker, 3D audio output, gesture recognition using VPL's DataGlove, a remote camera, and a BOOM-mounted CRT display.

The technologies from VIEW soon became commonplace in various military research organizations: US Air Force designed their Super Cockpit, British Aerospace their Virtual Cockpit, and the VECTA system for training pilots.

By the end of the 1980s, Virtual Reality research was booming everywhere. Unfortunately, most of it was either very expensive or specialized (or both), and the image quality was very low, which made the “reality” part of VR rather dubious. This did not stop the media, though, and the magazines raved about VR and how it was going to transform our lives. Of course, the technology could not live up to the expectations, and pretty soon Virtual Reality became just yet another buzzword. But research continued in silence.

Let us now see how a Virtual Reality experience can be achieved and then take a look at the current trends in VR technology.

1.2 Principles of immersivity

Murray [Murray 97] explains immersion as “a metaphorical term derived from the physical experience of being submerged in water”, and goes on to say that:

“...we seek the same feeling from a psychologically immersive experience: the sensation of being surrounded by a completely other reality, as different as water is from air, that takes over all of our attention, our whole perceptual apparatus.”

In Virtual Reality applications, immersion is usually defined as the full sensory replacement by artificial means, that is, all visual, audible, and tactile signals actually are computer-generated instead of originating in the real world. In VR applications, less developed senses such as taste and smell are usually ignored, as they are much more difficult to produce artificially than sights or sounds. Although we define our environment mostly by how it looks and sounds, the other senses are not really that important for a fully immersive experience [Pausch 97].

Immersion is mostly psychological: a good book or a film – not to mention the more engaging computer games – can cause temporary loss of knowledge of self or immediate surroundings. In fact, many current games offer better image and sound quality and a deeper sense of engagement and immersion than the VR applications a few years ago [Mapleson 94].

In this chapter, we will however not concentrate on the psychological side, but on the physiological side of what makes immersion real: how to produce physical images that are convincing enough.

1.2.1 Immersion and Senses

Sight is without a doubt our most important sensory channel. It has the ability to process more information than our other senses combined, and our brain can discern very small details very rapidly from the vast amount of visual data we see. A lot of the processing is already done before the image reaches consciousness, and we receive more than one kind of information about what we see: for example, the position of the eyes [Nienstedt 93].

The human eye is a wonderful instrument, and because of that, creating fully realistic images in real time is simply not possible at the moment. In order to have an immersive experience, the user needs to use *suspension of disbelief*, the ability to forget about the actual errors and limitations he perceives in the virtual world. Often, a very basic level of rendering is enough for an immersive experience, and other things, such as behavior and sound may contribute to the suspension of disbelief.

Sound is also important to an immersive system, however, satisfying audio experience can often be achieved using relatively easy means. Because the human directional hearing is not very accurate without the visual component (on the average 3.6 degrees in the frontal direction [Huopaniemi 99]), fooling the ear is much easier than the eye. Sound and its reproduction is discussed more in Chapter 2.7: *Audio Systems* on page 61.

The addition of *smell* and *taste* often diminishes the immersive experience at the current level of technology. Both senses are not understood nearly well enough to provide any sort of immersive feeling, and my personal experience confirms this. The whiff of a woman's perfume in the middle of a movie only reminds me of the actual theatre I am in, and does not add to the story. In the future this is hopefully remedied, but the technical implementation of "taste" emitters is rather elusive. Research on olfactory interfaces continues [Youngblut 96], [Göbel 00].

Touch and the *kinesthetic sense* are definitely the Holy Grails of VR. Currently, in order to provide tactile feedback, very complex machinery has to be installed, with multiple joints, servo motors, lots of wiring and the constant danger of electrocution. For one finger, this may not be such a problem, but a full hand/arm, not to mention the body, the result often is not satisfactory or even usable [Burdea 96]. Wearable solutions have begun to recently emerge, such as the Wearable Master [Iwata 99].

The skin and muscles are also very delicate and able to detect extremely small detail and vibrations. While their resolution and dynamics are not as good as the eyes', the sheer surface of the skin makes it difficult to make a system that would provide the sensation of touch for the whole body.

Using the current level of technology, images and sound are the easiest to produce, and they are quite enough to acquire "spectator" -level immersion. In the U.S. Department of Defense simulator program SIMNET the 60% rule was introduced and successfully used: "Attaining a 60% of reality was a high enough

level of fidelity to cause most simulation participants to become sufficiently immersed in the simulation experience to ignore the discrepancies of the missing 40%” [Neyland 97].

Let us now delve into the details of how the eye works, and how it can be fooled.

1.2.2 The Difficulty of Fooling The Eye

Accuracy, Accommodation, and Adaptation

The eye is extremely accurate and is able to discern details that are approximately 1 minute of an arc (1/60th of a degree) apart in ideal lighting conditions. This means that, for example, at the distance of 1m, the eye can separate elements that are in the order of 0.3 mm, provided that the contrast between the elements is optimal. In practice, however, the images produced by computer technology have far worse resolution, and this is usually compensated for by using different *antialiasing* techniques.

Accommodation means the capability of the eye to accommodate to different distances. Since the eye is not an ideal optical system (a so-called pinhole camera), only a certain part of the area we see depthwise is in clear focus at a time. This effect is called Depth of Field. On the average, the eye can accommodate to a distance between a few centimeters to 6 m, after which the eye is already accommodated to an infinite distance. The shortest accommodation distance increases with age, by approximately 1 cm/year.

The human eye is able to adapt to very different lighting conditions over a dynamic range of 10^{13} , and still see subtle differences. No man-made sensor can adapt to the same dynamic range, though at the both lower and upper ends certain type sensors do have better accuracy.

Field Of Vision

Figure 1, next page, shows the approximate field of vision of the human eyes. The overlapping area between the left and the right eye is the area where stereo vision functions. As can be seen, it is approximately 120 degrees wide, with the entire field of vision ranging to 200 degrees. This, combined with the accuracy of the eye (see below) makes it very difficult to produce artificial images that fool the eye to believe the image is real.

For example, an average person might be able to read a 10 pt. font at the distance of 1.5 m in average indoor lighting conditions - which makes the height of a letter approximately 5 arc minutes. In order to produce a computer display that displays the same size pixels at the approximately 35 cm distance to the screen, the character size may be 0.45 mm. If we assume it takes at least 8 pixels square to make a character, the resolution of the display must be approximately 450 DPI, a density not feasible at current technology for large displays¹.




Figure 1: The field of vision of the human (from [Kalawsky 93]).

Perception of Depth

There are two categories of depth cues: monocular and stereoscopic. Monocular cues are those that can be perceived through a single eye (or camera) and provide most of the depth cues. The monocular cues are according to [Lipton 97] with a deeper explanation in [Kaufman 74]:

- **Light and shade.** Objects can be made to look solid, or rounded, or to rest on surfaces by simply lighting and shading the object appropriately.
- **Relative size** tells us often the distance of an object: we know how big a car is, so we can judge its distance by its apparent size.
- **Interposition or occlusion.** Often overlooked, this simply means that things that are in the foreground cover objects that are in the background.
- **Textural gradient.** This is a modern observation, and simply means that the texture of an object is more clearly apparent when the object is closer to the observer.
- **Aerial perspective** is the diminution in visibility of distant objects caused by intervening haze or fog.

-
1. Small displays can achieve this and more when they are made using same techniques as are used for microelectronic circuits. However, the manufacturing technologies are not usable for large displays.

- **Motion parallax.** For example, watching landscape go by from a car: the distant hills move slower than telephone poles on the edge of the road.
- **Perspective,** sometimes called “geometric,” “rectilinear,” or “photographic perspective is the most important monocular depth cue. It means the relationship between foreground and background objects. A typical example of perspective is a pair of rails, that (on a plain) seem to converge at the horizon.
- **Depth cuing.** This is a common technique where you vary the brightness of an object depending on the distance to the viewer. Also known as “exaggerated aerial perspective”.

The monocular depth cues can provide by themselves a very three-dimensional image, like in everyday TV shows, or first-person 3D computer games. One of the most famous modern masters of playing with the eye is M.C. Escher, whose drawings produce stunning optical illusions of depth and impossible objects [Escher 88].

Stereopsis

It wasn't until 1838 when Wheatstone explained the mystery of our two-eye vision system and why we do not see two separate images. He introduced the concept of *stereopsis*, or the sense of depth by comparing two separate images [Wheatstone 38], explaining that the brain fuses the two images together, using this information to assess depth.

The first key concept in stereopsis is the *retinal disparity*, which means that if you could take pictures of the images on the retina and superimpose them, you would notice that the images are not exactly in the same position (assuming the eyes were not converged to look at the object itself).

The second key concept is *parallax*, which is what produces disparity. Parallax is caused by the fact that our eyes are not in the same position, separated by a distance that is commonly referred to as *interaxial separation*, or *interpupillary distance* (IPD).

Parallax may be positive, negative, or divergent. When the eyes are observing the images on the screen with zero parallax, the optical axes of the eyes cross at the projection plane. When image points have zero parallax, they are said to have *zero parallax setting* (ZPS). See Figure 2, below, for an illustration.

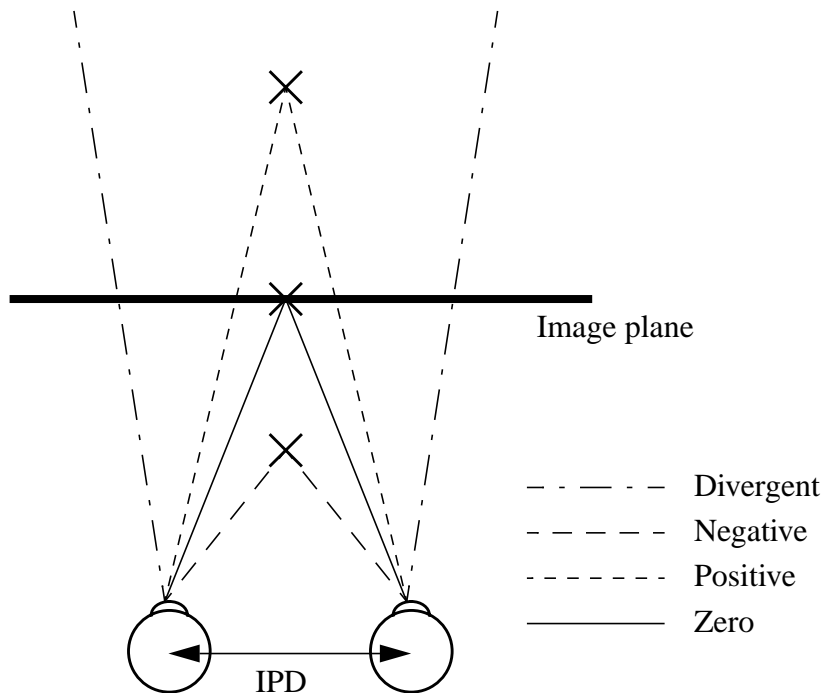


Figure 2: Parallax classifications and IPD.

1.3 How Three-Dimensional Displays Are Made?

There are several ways how the 3D displays have been attempted in the recent years. The first were probably the stereograms, which were made after the invention of photography, using a special camera that would take two pictures about 10 cm apart, and then the special pictures would be viewed through a simple device that would show the eyes the correct picture – the left picture to the left eye and vice versa. This tradition still lives on, with the ViewMaster toy from Fisher-Price [Sell 95].

For interactive computer graphics, we find that the basic technique used in the ViewMaster toy is still quite relevant. But why can't we just produce immersive images without the glasses, since they are rather cumbersome?

1.3.1 Autostereoscopic displays

Autostereoscopic displays give the viewer a 3D image without the aid of any auxiliary devices, such as glasses or headsets. Autostereoscopic displays are usually divided into three classes: re-imaging displays, volumetric displays, and parallax displays [Halle 97].

Re-imaging displays

Re-imaging displays do not produce a three-dimensional image by themselves, but they are used to affect the appearance of a 3D display in some way. The two simplest re-imaging displays are just plain glass and a mirror, since they do, in a sense, reconstruct an existing image. Using these simple devices or more complex optics, even fully three-dimensional displays can be created, and these have been used in games and parlor tricks for quite some time already [Kawamoto 95]. However, it is easier to construct generic synthetic 3D images with the other two techniques.

Volumetric displays

Volumetric displays work by filling or sweeping out a volume of space, using techniques such as a varifocal mirror display [Traub 67], or using some sort of medium that can be excited to emit light at certain points [Lewis 71]. What is good about the volumetric displays is that regardless of the technology used, they allow the viewer to look at the image from multiple locations, and even provide ocular accommodation (that is, there is actually something for the eyes to focus on).

Parallax displays

Parallax displays have a surface covered with display elements that can emit light in varying intensities to various directions. The commonly seen hologram is of this class, as are parallax barrier displays, lenticular sheet displays, time-multiplexed displays [Moore 92], etc. It is possible to also create the displays electronically, and compute the fringe patterns so that 3D images and movies can be shown within a “window” [Lucente 95].

The main problem with all three techniques is that they are all either hideously expensive, or unsuitable for interactive graphics, or both. Cheap, parallax displays can already be made (for example, like those in credit cards), but they are totally unsuitable for any scene containing movement. The most promising technique seems to be a parallax display that is generated electrically, and can thus be computer controlled. However, the technology is currently slow and very expensive, and does not produce very good quality images, making it not suitable for mass production.

1.3.2 Stereograms and Head-Mounted Displays

When true 3D is not available, that is, when we have to con the brain somehow into believing it is seeing a two-dimensional image, we go back to the ViewMaster-technology and start to look at how the human eye and brain function and how they discern depth. After all, the projections on the retina are two-dimensional, so in theory there should be no difference between artificially created images and real-world 3D objects.

To achieve stereopsis, the eyes must see a slightly different image (see *Stereopsis* on page 7), which can be acquired either by photographing (stereogram) or by computation. The resulting image pair must then be displayed to the viewer so that the eyes view their respective images, but do not see the image for the other eye. Stereopsis is lost, if this *crosstalk*, or *bleeding* is too severe.

The most straightforward way of displaying the left-right image pair is by having two separate optical systems, one for each eye, like in the ViewMaster or the old stereograms. This approach is generally used in Head-Mounted Displays (HMD), where the optical projection system is attached to the user's head. For more discussion about the current status of the HMD displays, please see *1.4.2 HMD devices* on page 13.

1.3.3 Stereographic Displays

When we have only a single display device available, or do not wish to encumber the user with a special display (like with printing press), it is possible to use normal computer monitors, screens, or paper. The following discusses how a single device can achieve the stereographic effect.

Image separation

With a single display we still need to separate the left-right eye images somehow. Typically the stereographic display displays alternating images for the eyes in a successive fashion, and the images are then separated using either active eyewear (LCD goggles, such as the StereoGraphics CrystalEyes [Lipton 90]) or passive eyewear (polarization glasses, red-green, or red-blue glasses [Haggerty 90]). In general, the passive eyewear tends to be a lot cheaper, and is thus used in large display environments such as amusement parks, whereas the quality of the more expensive active systems is better, and thus used in more serious applications.

Interlace Stereo

The original stereo-vision television format used to encode the left and right eye images on the odd and even fields of the normal PAL or NTSC image. The television signal is encoded so that the odd field contains all the odd lines and the even field contains all the even lines of the picture, and these are alternated at 60 Hz, resulting in a full field rate of 30 Hz (for PAL, 50 and 25 Hz, respectively). Unfortunately, the result is often flickery and has its vertical resolution cut in half (to approximately 240 lines for NTSC and 280 lines for PAL).

Field Sequential Stereo

For progressive² displays such as the current computer monitor displays, field sequential stereo (FSS) is a natural extension of interlace stereo. In FSS you send the left and right eye images separately on alternating frames, which preserves vertical resolution, but does obviously cut the refresh rate to half. This is the rea-

2. Non-interlaced - all fields are full resolution and correspond to one frame.

son this format hasn't been popular until the past few years, because in order to achieve flicker-free displays (at least 50 Hz for one eye) the display device must manage refresh rates in excess of 100 Hz.

Above-and-Below Stereo

The left and right eye fields can also be arranged to appear in a single frame on top of each other. When played back on a monitor which has double the field rate of the playback device, the images then appear as alternating images. Of course, vertical resolution is lost again in this format, but it is an easy method that works across a variety of interlaced and progressive storage and display devices.

There also exists a variant called side-by-side stereo, where the images are squeezed horizontally instead of vertically. This has the added bonus of maintaining the vertical resolution at the expense of horizontal resolution, which often leads to a smaller loss of image quality than with above-and-below format. Of course, the equipment needed to display such imagery is slightly more complex than with the simple above-and-below format.

White-Line-Code

The White-Line-Code (WLC) was developed for high-quality, but low-cost solution for consumer displays. It does not care whether the image is progressive, interlaced, above-and-below, but it uses a few pixels wide bar at the bottom of the screen, which, when 25% white and 75% black, signifies a left field and when 75% white and 25% black, signifies a right field.

Interleaved Stereo

VRex company has introduced a method where you have a special filter in front of the projector lens, that polarizes the even lines horizontally and odd lines vertically (or vice versa). The projection screen is made of a special material that preserves polarization and the image separation is handled through a pair of light glasses that have polarizing lenses [VRex 98].

This method obviously cuts the available vertical resolution in half (to approximately 400 lines) since the left and right eye images are interleaved on the even and odd scanlines.

1.3.4 It Ain't Real, Though...

The biggest difference between the virtual image created using methods detailed in previous chapters and the real world is that while eyes converge correctly on both, the eyes always accommodate to the actual object. In real-world this means the object itself, but with virtual image this means the projection plane. This problem is more noticeable the closer the projection plane is, and is considerably less with large projection screens that are far away (remember, the eye accommodates only to approximately 6 meters, see *Accuracy, Accommodation, and Adaptation* on page 5). This does not seem to be a very bad problem, though, and

especially children have no difficulty adjusting themselves [Lipton 97]. Inoue also finds that at longer distances the effect grows weaker, though the far-near accommodation time is longer after viewing stereo images [Inoue 90]. Using collimated displays allows the eye to accommodate at longer ranges by providing parallel light rays [GMO 99].

Also, in computer graphics the images are often calculated using a pinhole camera model, meaning that every object is in clear focus, which is clearly not the case with the eye. However, rendering this Depth-of-Field (DoF) effect is a multi-pass operation [Blythe 98], which impacts heavily on performance when the object is to produce real-time graphics³. Luckily the area of sharp vision is very small, and the users don't pay much attention to the lack of DoF-effects. In fact, since it is almost impossible to accurately gauge the direction of the user's attention, it is better that all areas are in perfect focus.

1.4 Current immersive VR display technology

1.4.1 Classification of VR display systems

Immersive VR displays can be divided into three distinct classes [Lantz 96], [Takala 96]:

1. Spatially Immersive Display (SID)
2. Virtual Model Display (VMD)
3. Head-Mounted Display (HMD)

They were all designed to tackle different areas of VR, and approach the problem slightly differently. HMD displays are the oldest display type, simply because the technology for stereo projection matured later than standard CRT projection. In fact, the earliest models of stereograms can be considered static versions of HMDs.

Standard desktop VR applications are sometimes applicable here, too. When stereo projection is used on a standard monitor, the result has many uses in different areas, such as virtual prototyping. Head tracking can also be added for greater feeling of depth [Deering 92]. These "aquarium" -type VR applications that are not immersive are not discussed here.

3. In movie industry the Depth-of-Field effect is, however, used with great success and heavily increases the authenticity of the pictures, since people are already accustomed to seeing such effects due to the imperfections in camera systems. It also serves as a means to focus the viewer's attention to whatever the director wants.

An excellent account on the currently existing commercial virtual environment display and interaction devices is available in [Youngblut 96]. This chapter will only take a short, general look at the different types of display systems currently in use.

1.4.2 HMD devices

The traditional image of Virtual Reality usually entails a helmet, a data suit or a data glove. In many cases, using a Head Mounted Display is a good idea, since they are relatively easy to build which makes them cheap. However, they do bring several problems with them:

- The resolution of the displays is generally poor. At the writing of this, the most advanced systems contain a resolution of 1280x1024 pixels, but they cost up to 60,000 USD. Cheaper versions go up to 640x480 pixels, which is not acceptable for an immersive experience.
- With current technology, low resolutions and close viewpoint cause the RGB triplets used to display colors to be visible. This causes breakup of the image colors, as the user is able to discern the different color elements instead of the intended, mixed color.
- The Field of View⁴ (FOV) of a typical HMD is very narrow, usually in the order of 60 degrees. There are helmets that can provide a FOV of up to 120 degrees or more, but then the low resolution becomes a problem, with pixels clearly visible to the user [Kalawsky 93].
- HMDs are usually very heavy, which incurs penalties for the user, if he has to wear them for prolonged periods. Typical problems include stiffness in the neck, headaches, and muscle tensions.
- The HMDs are vulnerable to latency problems, which in turn may cause nausea and disorientation in the users [Kalawsky 93].
- Head Mounted Displays do not accommodate multiple users very well: each has to have his own device to participate in the virtual experience.
- Psychological factors are an important limitation: engineers and decision-makers are reluctant to use such apparatus, considered game gear [Lantz 96].

The relative cheapness of HMDs has brought them to the consumer. Several units are already available, and the consumers are already available of Virtual Reality.

Figure 3 portrays a typical HMD display.

4. Field of View: the width and height of the displayable area measured in degrees.



Figure 3: Virtual Reality Head Mounted Display, which is freely worn.

BOOM -type displays

The HMD devices are usually categorized into two classes: Boom-devices, which have usually a larger screen, which is then suspended from a physical device such as a boom; and true head-mounted displays, where the full weight of the helmet is worn by the user.

Usually, the non-boom-operated displays offer greater flexibility and freedom to the user, but because the boom system may carry more weight, the image generation is better (more resolution, better field-of-view) on a boom-based system. In addition, booms can allow very precise measurement of the user's movements, and can in some cases even provide haptic feedback by using servo motors to prevent or hamper user's movements. Figure 4, below, displays a typical boom HMD.



Figure 4: The BOOM-3C device from FakeSpace, Inc.

1.4.3 Virtual Model Displays

A Virtual Model Display (VMD) is an extension of desktop VR, where the 3D model is viewed on a standard CRT screen. The screen size is enlarged to approximately 2-3 meters wide and 1.5-2.5 meters high, on which the image is

then projected in stereo. Usually these screens are made using back-projection CRT technology, and in many cases the technology is very similar to the CAVE (see *The CAVE Automatic Virtual Environment* on page 17).

Typically, a VMD is used to view CAD models, which thanks to the 3D-projection, appear to float on the screen. If a head-tracking system is used, the models can then be looked at from different viewpoints just by moving the head.

The VMD is finding many uses in the automobile and CAD industry, as well as scientific visualization [Obeysekare 96], since it allows a better view than the normal workstation screen and allows for multi-user collaboration. It is also considerably cheaper and easier to fit into a room than a CAVE. The main difference between a VMD and a Spatially Immersive Display (such as the CAVE) is the fact that the VMD is not immersive. The user will stay aware of the surrounding room and the display device itself, and a feeling of “being transported elsewhere” does not occur.

The ImmersaDesk

An ImmersaDesk is basically a huge monitor, which allows several people (usually two or three) to examine the same data. Unlike a conventional, 21” monitor, the ImmersaDesk is up to 2 meters in diameter, allowing a good view for all participants. See Figure 5, below.



Figure 5: The ImmersaDesk by Pyramid Systems.

In addition, the ImmersaDesk usually applies stereo vision in order to display 3D objects. It can also combine different manipulation tools, such as wands and cybergloves for enhanced user interaction. In most cases a head-tracking system is also available.

Even though the history of the ImmersaDesk (or Responsive Workbench [Krüger 94] or the Holobench, depending on the manufacturer) is similar to the CAVE (both were originally developed at the EVL), they do differ somewhat in their application, as the ImmersaDesk is definitely meant for non-immersive work [Pol 99]. The Holobench from TAN ProjectionsTechnologie (see Figure 6, below) is an attempt to bridge the gap between the CAVE and the ImmersaDesk, giving the user a more immersive environment.



Figure 6: The Holobench. Image courtesy of TAN GmbH.

1.4.4 Spatially Immersive Displays, such as CAVE

A Spatially Immersive Display (SID) extends the idea of an ImmersaDesk by surrounding the user with multiple projection screens, creating a much more effective and immersive experience. While the VMD is more like a piece of furniture, the SID will stay as much in the back as possible to help in the suspension of disbelief necessary when creating a convincing virtual environment. The SID can be considered successful, when the user is not aware of any display device at all.

One of the main problems with HMD displays is the lag they exhibit when the head is turned: the head is capable of making very rapid changes, and if the hardware and the software cannot follow fast enough, the viewer will see a noticeable lag. In some cases, this has been known to induce nausea [Burdea 96].

The Spatially Immersive Display fixes that by pre-calculating and pre-projecting the image to the different directions, and thus turning the head is a real-world operation which does not have to be translated to a virtual environment operation, saving considerable time and lag. It also removes the cumbersome display equipment from the encumbered user and makes it static. SIDs also scale to multiple users easily, and allow much larger models to be displayed at once than HMDs (where you have to keep turning your head because of the low FOV).

The obvious drawback to this approach is that the SIDs are invariably very large and often expensive. The display technology used requires much more complex systems, and image separation is not as easy as with HMD devices.

The CAVE Automatic Virtual Environment

A CAVE Automatic Virtual Environment (CAVE) is a cube-like structure, where 3-6 sides are used for back-projection. Typical CAVEs, such as the ones built and sold by Pyramid Systems, Inc.⁵ consist of 3 walls and a floor, which is the original design pioneered in the EVL laboratory in University of Illinois, Chicago in 1995 [Cruz-Neira 95]. A typical CAVE structure is shown in Figure 7.




Figure 7: A typical 4-sided CAVE structure. Notice the floor projection is from above.

The CAVE has been in many ways the first truly immersive experience. Since the image surrounds the user with only an approximately 90 degree angle missing directly to the back, the user is fully immersed into the virtual environment.

The CAVE has also been the first really successful VR device in the sense that while it is very expensive to build and maintain, it is actually useful, and can be worked with for long periods of time. It also allows the collaboration of several users, making it better suited for corporate needs. The CAVE has spawned a number of imitators, and a great many corporate research organizations and universities now have their own CAVE installation.

The CAVE consists usually of at least two walls and the floor (like the original EVL cave), the floor being included because it adds greatly to the degree of immersivity because the human FOV includes the ground in front of us (see Figure 1 on page 6. We also often look at the ground for reference and to avoid obstacles.) The newer CAVEs have usually at least three walls and the floor,

5. <http://www.pyramidsystems.com/>. They sell a commercialized version of the original CAVE from EVL.

while the more expensive ones have five or all six sides of the cube used. One such six-sided CAVE is at the Kungliga Tekniska Högskolan⁶, in Stockholm, Sweden (see Figure 8).




Figure 8: The KTH Six-sided CAVE built by TAN GmbH. Image courtesy of PDC/KTH.

Reality Center & IMAX

The Reality Center from Silicon Graphics, Inc. is a curved wall, where several projectors are used. It has some resemblance to multivision systems, but the images are joined to a single, large image. There are several similar systems available from a multitude of manufactures (see Table 1 on page 23).



Figure 9: The PRODAS curved screen from SEOS installed at SGI Reading, UK.

6. <http://www.pdc.kth.se/projects/vr-cube/>

Similar displays are also used in the IMAX theatres⁷, of which there are already several installed around the world. IMAX theatres use a very large screen (designed to fill most of the viewer's field of view) on which special projector (or projectors in case of 3D images) display a very high-resolution image from 70 mm film.

The IMAX theatres can also be built to show 3D images, using either standard shutter glasses or polarization goggles.

Domes

The dome is basically a CAVE, but with a curved overhead screen, much like in a planetarium. They are most often used in different motion rides and large theatres, such as the Omnimax theatres, and also in military simulators. See Figure 10.




Figure 10: The VisionDome from ARC. Image courtesy of Alternate Realities Corporation.

The dome systems usually produce a better-quality image than cubical shaped screens because of the equidistant projection surfaces, but on the other hand projections for curved screens become rather complex, especially when edge matching and maintaining standard luminosity over the whole projection area [Clapp 87]. Also, the geometry calculation may be a problem if the projection geometry is not fixed (for example, if the user moves and his location is tracked), increasing computational complexity. The construction of a dome is more complex than a CAVE, and does require more space. They are thus often expensive systems, but do produce very spectacular imagery.

Domes have also been built into other shapes, like dodecahedrons [Thomas 91], for which geometrical calculations are often easier and faster.

7. <http://www.imax.com/>. See also OMNIMAX.

1.4.5 Futuristic displays

The physical structure of the virtual display, such as the CAVE, prevents the users from moving too far in the virtual world, because sooner or later they will bump to the wall. This is why movement in the virtual environment needs to be done by using joysticks, 3D mice, and other input devices to move around in the virtual world and limit physical interaction to the minimum.

Even the HMDs are not immune to this problem, even though in theory they are not limited by the projection surfaces. In practice, they are limited by the length of the cables and physical room area. In fact, the space limitation is somewhat worse for HMDs, as the user cannot see his real, physical surroundings and is in danger of colliding with the physical objects within the same area.

Some solutions have been suggested like omni-directional treadmills [Darken 97], [Christensen 98], but none as interesting as the CyberSphere.

The CyberSphere

The scientists in VR-Systems UK have been researching a CyberSphere [Eyre 98], [Eureka 98], a device, which consists of a large, translucent sphere containing the user. The images are distortion-corrected and then projected on the surface of the sphere, allowing the user a full 360 degree field of view. It also allows the user to move around in the world, by walking inside the ball, which will move in response to the users movements. See Figure 11 for a drawing of the CyberSphere system.



Figure 11: The CyberSphere system. Copyright Volker Steiger/Science Photo Library

The advantages are obvious: this does allow a relatively natural interface to the virtual world, and will certainly be an impressive experience. However, the entire construction is very complex, especially the projection side. The Cyber-Sphere allows only one person inside, since the sphere can only be comfortable for one person. There is also no way of bringing extra equipment inside, as everything must be worn by the user and function without wiring. Also, the question of audio has not been addressed anywhere: the air cushion (800 Pa) that keeps the sphere afloat and as frictionless as possible is not very quiet, either. Virtual sound is supplied to the user via headphones [Eyre 00].

The angular momentum of the sphere is not zero, and will cause a noticeable effect. Julian Eyre writes [Eyre 00]:

“We don’t want the momentum of the sphere to be zero. Ideally it should be equal to the momentum of the occupant that they would have if moving in the real world. The sphere actually presents approximately twice the inertia of an average adult. We are therefore planning to introduce a form of active drive/braking control to the sphere which will reduce the perceived inertia to the correct level. In this way we will also be able to simulate slopes in the virtual environment.”

True 3D displays

It may well be that the autostereoscopic displays detailed in *1.3.1 Autostereoscopic displays* on page 8 are a part of the future. Currently the biggest problems seem to be the incredibly difficult computations required for producing electronic fringe patterns for parallax displays, as well as the cost of the necessary equipment to produce parallax or volumetric images. However, when considering Moore’s law and the unwritten rule that it takes approximately 20 years for a technology to mature from invention to popular consumption, it may well be that affordable 3D displays that no longer need any worn equipment will be available around 2010-2015. Some interesting progress has been made lately [Trayner 97], [Levin 99].

The Holodeck

In the 1987 science fiction series “Star Trek - The Next Generation”, the heroes use often a device called Holodeck, portrayed in Figure 12. It is a room, where hidden holographic and force field generators create a physical illusion of reality: The user can touch objects, break them, and even eat them. The computer can create intelligent agents, which in all respects look and feel like a fellow human being, giving an illusion of true intelligence.

The Holodeck is the true culmination of virtual reality. The user can physically act with his surroundings and completely feel that he is in a different world, where anything can happen - a real virtual world.




Figure 12: The Holodeck from Star Trek - The Next Generation. Star Trek is a registered trademark of Paramount Pictures.

Unfortunately, this device is still science fiction, but large virtual display devices are being constructed even today. And technology produces new innovations every day. We can already produce images and sound that come seemingly out of thin air [Halle 97], [Pompei 99], so the future may not be that far away.

2. CAVE Automatic Virtual Environment

In this chapter we will take a closer look at the CAVE, its construction, and the details of the construction. Since the CAVE is for the most part a pretty standard back-projection system that is just used for stereographic projection, what is stated here can be rather easily generalized for any configuration – not just the standard 3 side screens and a floor – CAVE.

CAVE Availability

It is not necessary to start building your own CAVE, since there already are several vendors of CAVE-like systems. The following table is an attempt to catalogue the known system providers. Most of these companies also supply ImmersaDesk-type displays, but they have been omitted for brevity. HMD vendors have also been omitted, simply because there are already several⁸. In addition to these companies that provide turnkey solutions, there are plenty of vendors that are capable of building a similar system, if specifications are given.

Table 1: Spatially Immersive Display vendors.

Vendor	Products
Pyramid Systems ^a	CAVE
FakeSpace ^b	RAVE, Immersive WorkWall, Immersive WorkRoom
MechDyne ^c	MD_Flex, LSVR, SSVR ^d
SEOS Displays Ltd. ^e	LUCID, MIDAS, PRODAS, PANORAMA
TAN GmbH ^f	TAN CUBE, TAN HoloBench, Responsive Workbench, TAN StereoVision, TAN PANORAMA
ProSolvias ^g	3D-CUBE
G.R.A.F ^h	ImmerZimmer, 3D Theatre
Alternate Realities Corporation ⁱ	VisionDome
CGSD Corporation ^j	Custom systems
Evans&Sutherland ^k	VistaView
Flogiston ^l	flostation
Goto Optical Mfg. Co.	Virtuarium
Spitz, Inc. ^m	ElectricSky, ElectricHorizon

8. For a good list of HMD vendors, please see the Visual Displays FAQ at <http://www.hitl.washington.edu/scivw/visual-faq.html>

Table 1: Spatially Immersive Display vendors.

Vendor	Products
VRex ⁿ	VR COVE, iPod
Sogitec Industrie ^o	Custom systems
Trimension, Inc. ^p	ReaCTor, V-Desk
BARCO ^q	Baron, Consul
SGI ^r	RealityCenter
HP ^s	Visualize
Panoram ^t	GVR-120, PanoWall, ViewStation
Concurrent Technologies ^u	CTC VR Bench, UIC CAVE

- a. Was merged with FakeSpace August 9, 1999. <http://www.pyramidsystems.com/>
- b. <http://www.fakespace.com/>
- c. <http://www.mechdyne.com/>
- d. A PC-based version is also available.
- e. <http://www.seos.co.uk/>
- f. <http://www.tan.de/>
- g. Filed for bankruptcy in 1999.
- h. <http://www.graf-factory.se/>
- i. <http://www.virtual-reality.com/>
- j. <http://www.cgsd.com/>
- k. <http://www.es.com/>
- l. <http://www.flogiston.com/>
- m. Specialized in planetarium systems. <http://www.spitzinc.com/>
- n. <http://www.vrex.com/>
- o. French company, <http://www.sogitec.com/>
- p. <http://www.trimension-inc.com/>
- q. <http://www.barco.com/>. VMDs and C³I-displays only.
- r. <http://www.sgi.com/>. The RealityCenter is a family of products, under which several Pyramid, SEOS, Trimension, and other products are sold.
- s. <http://www.hp.com/visualize/products/immersive/index.html>. Systems are provided by Panoram.
- t. <http://www.panoramtech.com/>
- u. <http://www.vr.ctc.com/>

It should be noted that the computer industry moves so fast that any paper document describing the current situation is out of date the day it is published, so it is suggested that the reader refers to the Visual Displays FAQ or some other internet resource for a more up-to-date list. The list given here is for illustrating the fact that within the past two years, immersive VR has become rather commonplace and is readily available, though still rather expensive. A typical turn-key solution for a 4-sided CAVE would be around 500,000 €, without the computing equipment, which would cost approximately 300,000 - 500,000 € more, depending on the desired image quality.

2.1 Physical Structure

The minimum CAVE dimensions, 3x3x3 meters are constrained by the human physiology. If the user stands in the middle of the CAVE, the ceiling structures should not be visible to him, which, according to Figure 1, means that the top of the projection screens should be above approximately 45 degrees in order to fall out of the users Field of View. This means that, assuming a 170 cm tall observer standing approximately 150 cm from the screen should have a 320 cm tall screen (150+170 cm). Luckily, since the human eye is not very good on detail at the outer edges of the vision, we can do with slightly less, especially if the ceiling does not have any distinguishing features and is very dark in comparison with the rest of the CAVE. In practice, the screens should be at least 250 cm high in order not to disturb the immersive feeling.

With the physical dimensions of the CAVE established, we can easily calculate the required distance for the projectors. Since we are using an image that has an aspect ratio of 1, that is, square, and practically all CRT tubes have an aspect ratio of 4/3 (which is the standard TV aspect ratio), we need to calculate the distance for a 4 meters wide and 3 meters high picture.

The projection distance is easily available from different projector manufacturers. For example, for the BARCO 1209s projector the distance is given in [Barco 97b], and states that for a 4 m wide image the physical distance (or the focal length) should be 5.28 m between the projector and the screen. See Figure 13, below for an illustration of the different distances.

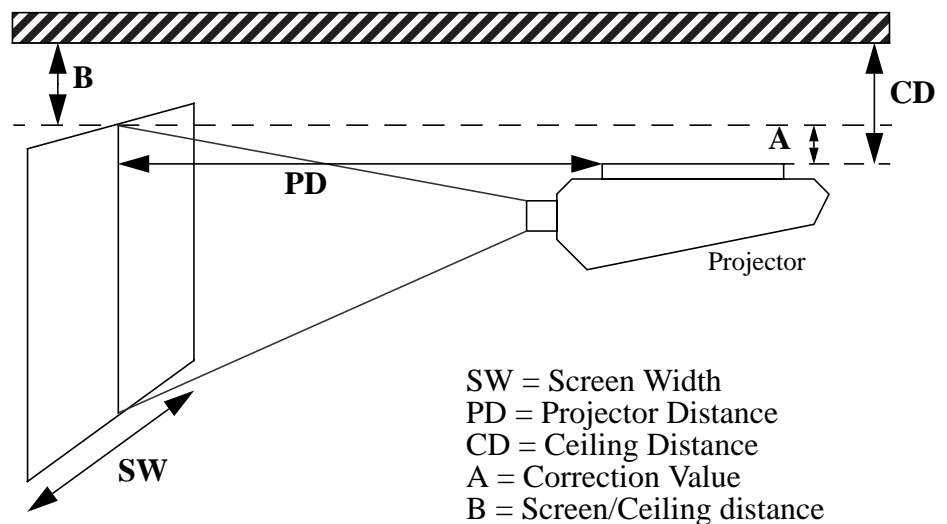


Figure 13: Projector distances.

In practice it is very difficult to find a room that is 6 m (5.28+projector size) + 3 m + 6 m = 15 m square, where most of the space is wasted. The required space can however be lessened considerably by using mirrors to cut down required physical distance, referred to as “folding the optics”. Figure 14 displays one possible configuration for the CAVE projectors and screen.




Figure 14: CAVE layout, top view. Projectors on the sides.

2.1.1 Mirror positioning

With the mirrors reflecting the image and shortening the distance the situation is complicated somewhat. The setup with one mirror is displayed in Figure 15.

The distance for the mirrors may be calculated using Equation 2–1:

$$D = F - d_c + \frac{d_c \sin \beta \sin \alpha}{\sin(\alpha + \beta + \gamma)}, \quad (2-1)$$

where D = distance required between screen and wall to fit a mirror, F is the focal length of the projector, d_c is the distance from the focal point of the projector along the optical axis to the screen, α is the angle between the mirror and the screen, β is the half-angle of the projector’s display (easily obtainable by using formula $\tan \beta = f/2w$, where f = focal length and w = image width), and γ is the angle between the projector and the screen. Note that this does not take into account the shear that happens if the projector is not located vertically at the half-way of the screen. However, this shear may be calculated easily, please see *Appendix A: Projector Geometry* on page 107.

Some other possible configurations are shown in Figure 16, displaying both one-mirror and two-mirror solutions. This configuration has the advantage that the image shear is easier to calculate and with the two-mirror solution the free space




Figure 15: Single-mirror configuration and parameters.

required is less than with a one-mirror solution. Also dust may be less problematic if the reflecting surface of the mirror points down. The downside is of course the increased cost and complexity, as well as the extra reflection which loses light. See *Mirrors* on page 42 for more discussion about the general properties of mirrors.




Figure 16: Alternate configurations for the Cave side projectors.

The size of the required mirror is of course dependent on the distance between the mirror and the projector. Bring the mirror closer, and you need a smaller mirror; bring it farther away and the mirror size grows. Larger mirrors tend to be quite a lot more expensive than smaller mirrors (remember, when distance is doubled, the area is quadrupled), but on the other hand, possible surface irregularities will have a bigger effect on the projected image if the mirror is closer to the projector. Also, larger mirrors are more difficult to make, and are more probable to have surface irregularities.

There really is no hard and fast answer on the question of the mirror placement. It is a complex optimization task, where the main problem is balancing the image quality, cost, and available space.

2.1.2 Height & Floor Projection

The height of the required installation depends greatly on the actual configuration of the Cave. If a full floor and ceiling projection is desired, the Cave must be lifted up from the ground, and a back-projection solution must be used for all sides of the cube. If a four-sided Cave is made, the height can be reduced significantly by providing a front-screen projection solution by projecting the image from the top to the floor using a mirror. See Figure 17 for the two different construction types. It is also possible to project the floor with the projector in the back of the Cave, which does change the construction somewhat.

Projecting the image to the floor with the projector on the top is of course the easier and simpler solution, but unfortunately the shadow of the user will be visible. The impact may be reduced somewhat by projecting the floor image slightly from the front of the CAVE, so that the users shadow falls behind him. In practice, this works fine.

The required height of the room may be calculated similarly to the required wall space, except that one must take the Cave structure into account so that no structural elements come into the light cone, which would provide disturbing shadows into the image. Neither should the mirror nor the projector be actually inside the Cave. Use Equation 2–1 on page 26 to calculate the height as well.

It is easy to see that the approximate height requirement for a 5-sided cave is at least $3+2+2 = 7$ meters of free space, which makes it unsuitable for many rooms. The minimum height of a 4-sided cave is approximately 3.8 meters.

2.1.3 Frames

The single most complex hardware system that needs to be built for a CAVE is the frame, which holds the screen in place. There are multiple configurations available from different manufacturers.




Figure 17: Cave layout, side view. Left side shows a typical CAVE with a top projection, the right side shows a CAVE with both floor and ceiling projections.

The original EVL CAVE [Cruz-Neira 95] used a simple steel and aluminum frame, which had a 9x3 meter screen stretched by 1/8" steel cable under tension. This construction, while simple, has four major drawbacks:

- Good, 9 meters long seamless back-projection screens are hard to find. There are some manufacturers, but with a smaller size you get a better selection of materials.
- The 1/8" inch cable (3.2 mm) is approximately the size of a single (physical) pixel. While this does not remove any parts of the image, it is extremely visible and reminds the user constantly about the real world, reducing the sense of immersion. In addition, this gives a very strong cue about the actual location of the walls.
- In order for the screens to be fully planar, the cable must be stretched very tight, causing extra problems with the construction of the rest of the frame.
- The steel cabling will cause disturbances in a magnetic tracker, whether it be an AC or DC tracker.

Most of these problems have been addressed in the TAN CUBE, which uses welded seams and wooden frames in stead of the steel cabling, resulting in a very satisfying immersive image. With proper projector calibration, any visible seams can be removed, and the user will not be able to discern the location or distance of the walls.

Material

Typical construction materials for frames are steel, aluminum, or wood. Steel, while otherwise easy to handle, is ferrous and thus will cause problems with most types of trackers (see 2.4.2 *Tracking Hardware* on page 51). Aluminum is light, versatile, and easy to work on, and will not cause too much problems with DC trackers, since it is not ferromagnetic. Wood is of course the ideal material considering magnetic tracking, since it is non-conducting. Wooden frames are usually bigger than metal frames, and require considerable expertise to build, as wood is a living material and may change shape if not treated correctly. Also care must be taken that not too much metal is used in its construction, since nails, screws, bolts, etc. may offset any benefit the wooden frame gives.

2.1.4 Fully enclosed CAVEs

A fully enclosed CAVE is one which has six sides, that is, it offers a 360 degree visual field. One obvious question is “how to get in”? The second question is “how to get all of the wiring inside?”

The most common solution is to make one wall of the CAVE hinged. On a system built with hard projection screens it is even possible to make only a part of the wall hinged (like a door), but usually the whole wall turns to avoid any extra seams, or hinges that might be visible to the user.

In some installations, such as the RAVE from FakeSpace⁹ the entire projection system is enclosed within the wall, and thus the projection system turns with the wall. This means that the projectors do not have to be recalibrated, but it does add to the complexity and weight of the turning system. Also, the projectors do not like even soft impacts, and lose their calibration easily.

Wiring is usually done by either leaving a small hole in one of the corners or by using wireless systems, such as the Ascension wireless trackers [Ascension 99].

Since the entire system is closed, air recycling is a problem, especially with multiple users. The VR-CUBE at KTH in Stockholm are planning an automatic ventilation system, that blows new air into the CAVE immediately when the door is opened, thus forcing the air to be recycled [Barth 00].

9. <http://www.fakespace.com/>

2.2 Back-projection systems

2.2.1 Image Separation

Stereographic projection in a CAVE environment is usually done using active eyewear, such as shutter glasses. While more expensive than passive eyewear, the shutter glasses are immune to user head rotation and positioning, whereas typical polarization-based solutions are susceptible to crosstalk: if the viewer turns his head even slightly to the side, the linear polarization filtering is no longer perfect. Using circular polarization this problem can be removed, but the sensitivity to incident light makes crosstalk still possible [Haggerty 90].

As mentioned in section 1.3.3 *Stereographic Displays* on page 10, there are both passive and active eyewear for image separation. On a back-projected systems, using passive eyewear such as polarizing glasses is difficult, because the plastic screen tends to remove polarization completely, since it attempts to be as diffusing as possible. TAN GmbH builds a metallic back-projection screen which preserves polarization, but it is very expensive.

While not suitable for a CAVE, polarization-based stereographic projection has a good quality/price ratio, making it a good choice for different Location-Based Entertainment (LBE) facilities.

2.2.2 Screen Type

Back projection is done on two different screen types: a soft *film screen*, and an hard *optical screen*. The soft screen is made out of soft plastic canvas, and the hard screen, a projection plate, is usually some sort of acrylic material [DNP]. The main differences between these screen types are:

- **Material:** The soft screens, such as those manufactured by Da-Lite, Stewart, or Harkness, are made out of elastic plastic approximately 0.3 mm thick, and need to be stretched over a frame in order to be smooth. The hard screens are made out of thicker acrylic, and are free-standing.
- **Weight:** Optical screens are heavier, whereas the film screens are considerably lighter.
- **Transportation:** Optical screens require much space, but the film screens can be folded into very small space.
- **Cost:** Large optical screens are very expensive, approximately three times the price of a top-of-the-line film screen. Film screens are available also in very low price, though their quality may leave something to be desired.
- **Diffusive qualities:** The material and thickness of hard screen allows manufacturers to place different optical qualities to the screen itself, such as Fresnel lenses. This allows the light rays to be made parallel, and considerably reduces hot spotting (see 2.2.4 *Hot Spots and Other Odd Problems* on page 34). Optical screens can also be made very resistant against ambient

light.

- **Viewing angle:** The film screens have a typical viewing angle around 70 degrees (see Figure 20 on page 35 for an illustration), whereas high-end optical screens have a viewing angle of up to 90 degrees, because of their better diffusive qualities. The gain also stays more constant across different viewing angles with some models of optical screens, but some do suffer from very bad ringing effect. Some optical screens can also be configured to project the light to a specific area: for example, to have a very wide viewing angle vertically, but a very narrow one horizontally. This has the effect of concentrating the light to the viewing area.
- **Gain:** The typical (on-axis) gain¹⁰ of a film screen is around 2.3, while the gain of a high-end optical screen is around 3.5, which makes them somewhat brighter. In general, higher gain screens provide a brighter image, but reduce the viewing angle.
- **Strength:** Since the soft screens are thin plastic, they are much more susceptible to wear and tear than hard screens.
- **Installation to a CAVE:** Film screens need a frame with elaborate schemes for edges, but optical screens can do with much less, since very thin edges can be achieved just by attaching the screens to each other.

Even though optical projection screens do produce a better quality image, their use has been limited in CAVE-like environments. This is probably because optical screen technology is still new, and rather expensive in larger configurations. The viewing angle dependency of cheaper optical screens is not very good for CAVE-applications, either. However, there is at least one manufacturer who specializes in optical screen-based CAVEs¹¹.

2.2.3 Brightness and Contrast

One of the problems with CRT systems is that they do not produce very bright images. Standard CRT projectors have brightnesses of approximately 200-250 ANSI lumens, whereas light-valve DMD projectors, for example, can achieve in the excess of 10,000 ANSI lumens [Barco 99], which is quite enough to be bright and clear in daylight. Normal CRT image “washes out” if the ambient light level is high enough. See 2.3 *Projectors* on page 44 for more information.

Using an active image separation system will also cause the apparent image brightness to be halved, since only one of the lenses is open at a given moment, reducing the light available to both eyes to half. Since the eye is not linear, the perceived brightness is not halved, but still visibly diminished.

10. Gain is the measure of apparent brightness compared to a standardized matte white surface (Gain=1.0).

11. Trimension, Inc. See *CAVE Availability* on page 23.

In order to achieve a brighter image a secondary projector may be added, that is either used to display the same image (so that the image brightness increases) or one projector can be used to display the left eye image while the other displays the right eye image. The former method is suitable for CRT projectors and shutter glasses, while the latter is more suitable with polarization-based screens. The increased cost and need of accuracy in placement often prevent this solution from being feasible, though.

With bright images the absorptive qualities of the screen become more important and light spillage from one screen to another screen will be a problem. The better the absorptive properties on the screen material has, the brighter images can be used [Cruz-Neira 95]. Figure 18, below, shows a case in the CAVE where the light spills from the front screen to the side screen and the floor (which have no image to give an illustration of the magnitude). While the eye will accommodate to the brighter surroundings, it will also see details better, and thus issues easily ignored in darker surroundings, such as cables, and seams, will be more visible.



Figure 18: Light spillage from one screen to the next.

Ambient light is a problem with back-projection systems, especially if the light has access to the projection area behind the screen. The effect of ambient light can be reduced by either increasing the projector brightness or by removing any ambient light. At least the area where the projector and optical system lies should be closed off to any ambient light, and also covered with black matte surface, which absorbs any ambient light from the projector itself.

Typical back-projection film screens have a transmittance of 40-50%, which means that half the light is transmitted and the other half either reflected back or absorbed. With such low transmittance, the CAVE room must be shielded against ambient light from the CAVE screens.

Figure 19, below, shows the path of the light to the viewer's eye through the different layers and how the light is absorbed. As can be seen, the intensity the viewer sees is only about 1/4 of the original image intensity.




Figure 19: The absorption of light.

2.2.4 Hot Spots and Other Odd Problems

Hot Spot

Hot spot is a name for a common occurrence in a back-projection system, when there seems to be a brighter spot in the image directly in the front of the viewer, which also seems to move with him [Kirkpatrick 91]. This happens because the diffusing characteristics of the screen are not ideal, and intensity is greater in the direction of the incident light. See Figure 20 for an illustration.

If the screen was ideal, the incoming ray would diffuse spherically to all directions and thus the brightness arriving at the viewer's eye would be constant across the whole screen area (assuming a constant image, such as a plain white field). Since the screen is not ideal, areas of the screen which lie directly between the projector and the eye will seem brighter. This is the reason why the hotspot seems to move with the user.

Better screens can reduce this effect considerably. Most back-projection plates contain a Fresnel lens, which turns the light rays perpendicular to the screen surface. This removes dependency on the incidence angle λ , and reduces hotspot considerably (but does not remove it, unless the screen is also an ideal diffuser). Some better film screens can also achieve a similar effect. See Figure 21, below.

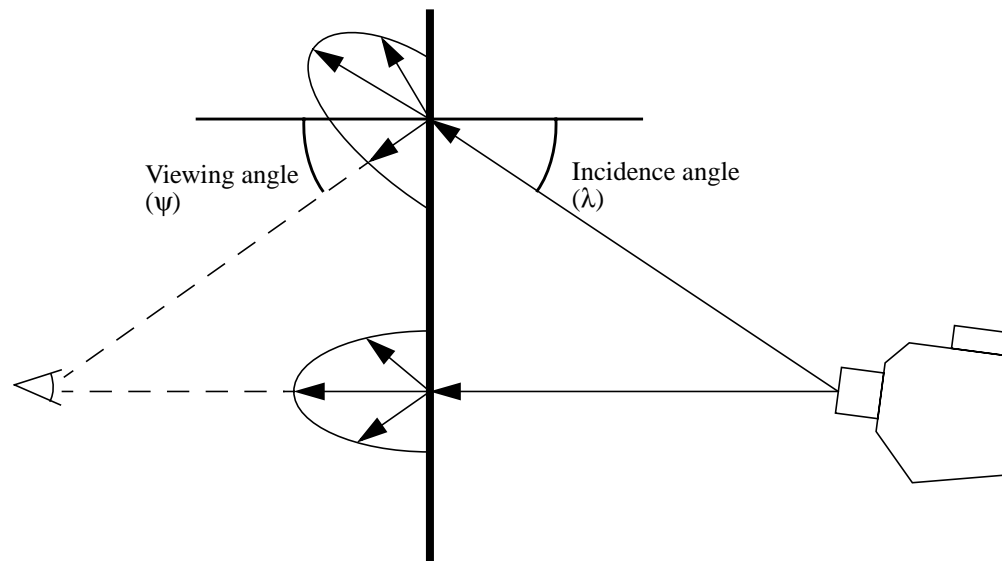


Figure 20: Hot spot effect and the projection geometry.




Figure 21: The Fresnel lens.

In order to achieve an even brightness display, the optical axis of the projector should cross the screen plane in an angle that keeps the incidence angle changes as small as possible, that is, the focal distance should be as long as possible, and the point of shortest distance between the screen and the projector should lie as close to the screen center as possible.

Hot spotting also occurs in front-projection systems [Kirkpatrick 91]. However, modern screens and projection technology no longer suffer seriously from hotspotting, and all major manufacturers can offer virtually hot-spot-free screens for both back and front projection. Front-projection hot spots can also be reduced by using a curved display.

While having good diffusive properties is important in eliminating the hot-spot, it will also cause smearing, reducing the overall image quality. For example individual pixels will be harder to distinguish. The overall screen brightness will also be reduced, because the light is spread more evenly across the whole hemisphere, instead of being concentrated towards the viewer (except when a Fresnel lens is used). In a CAVE-construction a lot of the diffused light will also hit the neighboring screens and be reflected to the viewer's eyes, causing uneven brightness distribution.

Distance variation

Obviously, the distance from the projectors to the different parts of the screen varies, which does cause some brightness loss (light intensity is reduced by $1/r^2$). On a three-meter screen this amounts to a maximum intensity difference of 2.2 (distance = 7.2 m) if the projector is mounted on the floor at 5.2 m from the screen. While most modern projectors are able to adjust to this (at least an option is usually available), it is usually a good idea to mount the projector so that the closest point to the projector is at the center of the screen.

As mentioned before, the screen is not an ideal diffuser and thus the gain is dependent on the viewing angle ψ . When the viewer moves in the CAVE, the angle between him and any point on the screen (i.e. the viewing angle) changes. If the surface brightness is also dependent on the incidence angle λ (as with a non-diffusive surface), the real *light angle* at which the light rays are seen is $\Phi = \psi + \lambda$. Table 2, below, shows the changes in the light angle with respect to the top and bottom of a 3 m screen as the user moves nearer and farther. The eye level is at 1.6 meters. This table is a rough estimate only, since it does not take into account that the distance to the corners of the screen is more than it is to the center of the edge (by a factor of $\sqrt{2}$).

Table 2: Viewing angle variation by distance.

Projector mount	View Direction	Light Angle (degrees)	
		Near (d=0.5m)	Far (d=1.5 m)
Low (floor level, d=0.3m)	Down	75	50
	Up	97	70
High (eye level, d=1.5m)	Down	88	63
	Up	86	59

It is easy to see from the table that it is rather difficult to get an even gain across the whole surface. However, the maximum allowable light angle of the screen should be at least 70 degrees, preferably more.

2.2.5 Floor

The floor is often a problematic issue, since it needs to withstand the weight of several users without breaking or bending significantly. Also, it may not introduce any disturbing audio elements (like creaking), and it must match the reflective properties and color of the wall.

Construction

The floor may be either projected from above (equivalent to front-projection) or below (back-projection). A back-projected floor must of course be translucent, and have the same transmissive properties as the back-projection screens on the wall. There are two alternatives:

- Use a translucent hard plastic (such as acrylic) surface. It may be difficult to find a provider that can build such a big surface that is thick and strong enough to support the required weight, but it is possible. For example, the TAN CUBE at Fraunhofer Institute IGD in Darmstadt has a 2.5x2.5 plastic floor, which is excellently matched against the color of the screens, making it very hard to discern any edges.
- Use a fully transparent surface, then cover it with the screen material. This is a slightly riskier, though easier method. The screen material, being plastic, is easily subject to wear and tear, which is why it must be fastened very tightly on the actual floor with something that is fully transparent as well. This approach will work well with both front- and back-projected screens. The six-sided VR-CUBE at KTH, Stockholm has been built using this method [Barth 00]. The canvas can also be protected by using a transparent film material, that can be replaced when it shows signs of wear.

For a front-projected screen the situation is a lot easier: since the floor may be opaque, normal construction methods and materials (such as wood) can be used. The original EVL CAVE used simple, painted wooden floor with quick color matching made in a local paint shop [Cruz-Neira 95]. More complex systems may also be devised, especially if things such as loudspeakers, vibrators, sensors, or cabling need to be placed underneath the floor.

One thing that should be remembered is dampening. The space between the building floor and the CAVE floor needs to be filled with material that absorbs sound, or else the acoustical properties of the CAVE will suffer.

It is also a good idea to leave room for later cabling, in case someone wishes to put sensors, or any other equipment under the floor. For example, in the Cyber-Stage built by TAN in Sankt Augustin, Germany, the floor has been equipped with low-frequency vibrators, that complement the audio system and increase the sense of immersion [Dai 97].

Surface

With the back-projected surface there isn't much choice in materials, since the translucency, transmission properties, and strength are the main constraints. However, care should be taken that the surface is not too slippery, or dangerous (and possibly costly) accidents may occur.

With the front-projected floor there are more possibilities:

- Take a normal wooden floor, then paint it with a color that matches the screen color. This is a simple method that usually produces satisfactory images, but searching the correct surface/paint combination is a problem. Commercial ventures have their own tried-and-true color matching methods, and produce good results.
- Plastic over wood: The floor can be covered with some suitable material, such as a hard plastic reflective screen. This is also a valid method, but may be more costly with no significant improvement in image quality.
- Screen material over wood. Finally, standard front-projection screen material may be stretched over the floor in order to get better reflecting properties. This is, of course, somewhat risky because the screen material may easily be damaged if mishandled. Also, film screens may move under the foot, making it dangerous. These problems can be levied somewhat by attaching the screen very tightly to the floor with double-sided tape and forbidding the use of shoes or other hard materials in the CAVE.

In every case, color matching will be a problem. It is very difficult to get the same combination of reflectance/transmittance for the floor and walls, and in most cases, perfect matching will fail. However, some of the color matching can also be made with the projectors, since the more expensive systems often are able to adjust the gain for each color separately, set the color temperature, and so on.

In the typical uses of the CAVE the difference between the walls and the floor is not critical. As long as the relative intensity of the walls and the floor matches, the floor image can get away with more errors, since it is not the main object of interest. Most of the user's attention will on the front and side screens.

2.2.6 Corners and Edges

Corners and edges are possibly the hardest problem in a CAVE after color matching to solve. The original EVL CAVE and its solution was already briefly touched in section 2.1.3 *Frames* on page 28, and elaborated here further. I will also examine some other alternative ways of making the corners.

The construction of the edges is important because that is the place where two images meet. If there is the slightest discontinuity in update rates, luminosity, calibration, or location, the eye will easily notice it. Maintaining these factors across one display area is relatively easy, as the rendering software, display hardware, and projectors mostly take care of such matters¹², but the edges are a problem to make fully continuous.

The edges need to be as continuous as possible, since occlusion is a stronger depth cue than stereopsis, and having an edge “cut away” a part of the image will make the system feel less immersive.

With hard projection surfaces cube edges are not such a big problem. Optical screens are free-standing and do not require big external structure to keep the under constant tension, unlike film screens. Edges are thus much easier to make seamless. If the film screen is not tightly fastened, its large surface will easily start to vibrate and move according to the air currents inside the CAVE room. Even people moving around near the screen will cause it to vibrate visibly. Since film screens are made of elastic material, it will take them some time to settle in, and will require later readjustment.

On a panoramic display there are more options to choose from: the pictures can just touch each other, or they can overlap slightly. Many commercial solutions exist for these kind of displays, for example [Mayer 96].

Let us now see a few methods for stretching the corners of a canvas screen. Note that these mostly apply to walls only. The floor and ceiling have to be made of a hard plastic, and are thus very easy to put into place with simple joints.

Steel Stretched Edges

The original EVL construction is shown in Figure 22, below. This construction has the advantage of relative simplicity, but if the tension on the cables is not high enough, the edges may “fall in” in an effect resembling a pincushion.

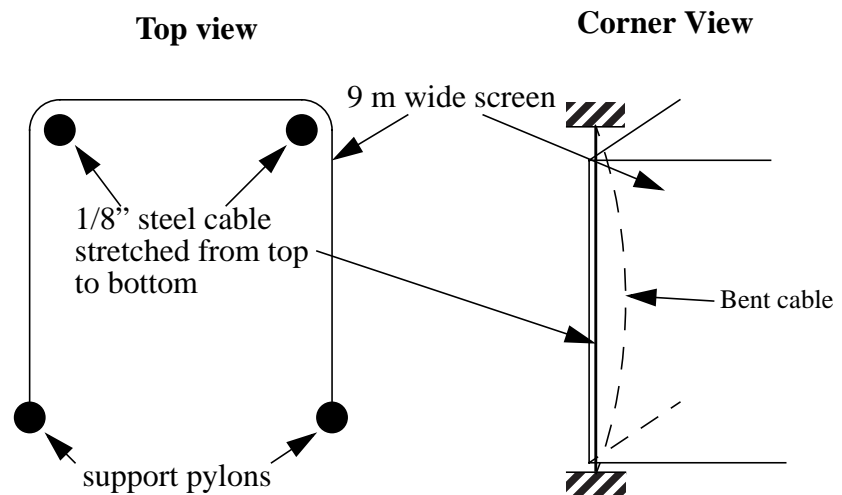


Figure 22: EVL/Pyramid Systems CAVE corner solution.

12. There will still be some artifacts left, see section 2.3 *Projectors* on page 44.

Cross Corner Edges

Another solution is presented in Figure 23. This time the screens have been interleaved and stretched separately. This has the advantage that screens are separate and easier to adjust, but cutting the screens has its risks and will reduce the overall strength of the construction. Of course, if one screen is faulty, it can be easily replaced without removing other screens.

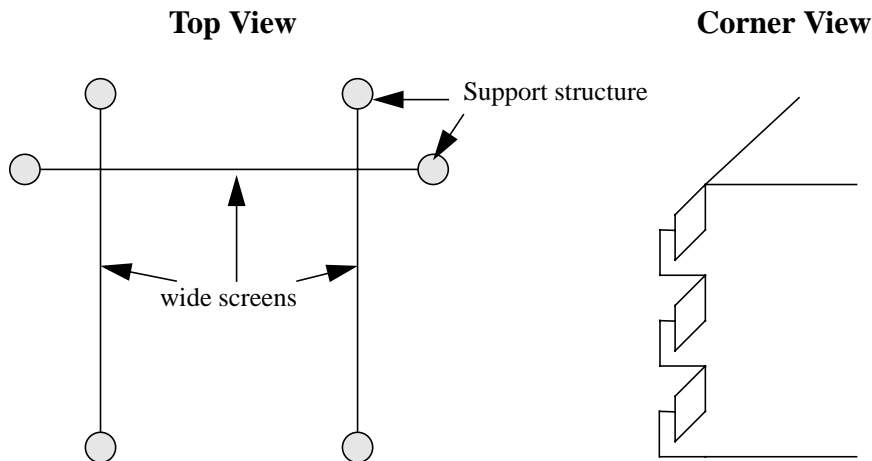


Figure 23: Cross Edge solution.

While the corners have to be carefully made, the end result is very good and produces no bad artifacts.

Welded Edges

One method that combines the advantages of the previous solutions is to seam (“weld”) the edges, which is displayed in Figure 24. This has the advantage of an even pull across the whole canvas, and almost no visible seam whatsoever. The solution is not very complex, either, but most of the stretching tension is now placed on the seam. Under normal use the seam will hold, but in case of an accidental collision by the user or an object to the walls, the seams may break. The whole canvas needs to be taken down and replaced, resulting in costly downtime.

Floor and Ceiling

For the floor there are several possibilities, and all are rather equal as long as the joint is even and does not introduce artifacts. One possible fastening method is shown in Figure 25. This is a standard, good, and solid method that keeps the canvas straight and produces no visible seam. The fastening plate reduces the possibility of tear. Unfortunately, this method does not allow later readjustments easily, so care must be taken during initial installation.

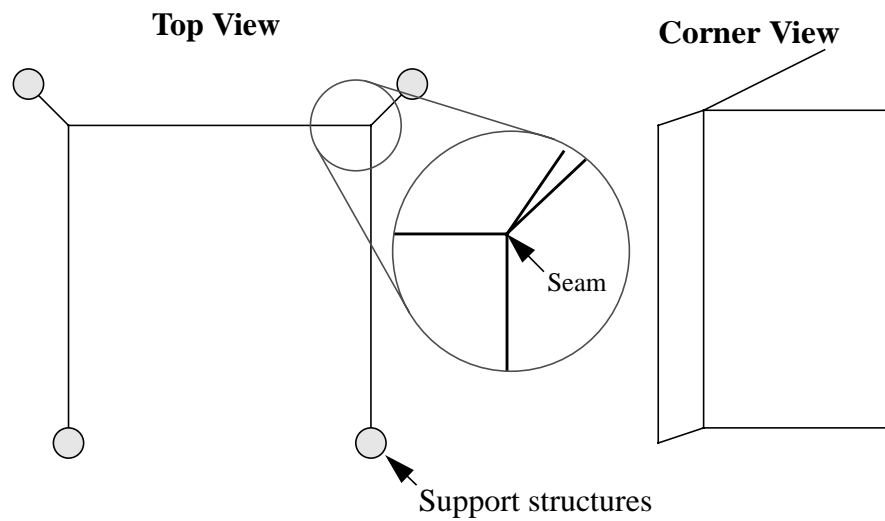


Figure 24: Welding or seaming the edges.




Figure 25: Joining the (front-projected) floor seamlessly with the walls.

Connection between the ceiling and the wall is similar to the wall-to-wall connection, or wall-to-floor connection, depending on whether the ceiling is a canvas or a projection plate.

Sometimes the ceiling is made of film screen, too, and a construction as in Figure 26 might be used.




Figure 26: Sharp edges

2.2.7 Mirrors

Mirrors

There are two types of mirrors that are useful for back-projection systems: foil mirrors and surface glass mirrors. Normal mirrors that have the reflective surface on the backside of the glass are not usually used, since they have several reflecting surfaces, which may produce double images in some cases.

Surface glass mirrors have their reflective layer topmost and thus only reflect a single image. Because of their construction, they are more susceptible to damage than normal mirrors.

Foil mirrors are usually made out of some reflective metallic foil (such as mylar) stretched over a frame. They are very light and as such, easy to handle, and they are not as expensive as full-scale glass mirrors, but they also tend to be more fragile, as the foil is usually very thin. Their reflectivity is usually less than for glass-mirrors, in the order of 85-95%, whereas the reflectivity of glass mirrors is usually between 90-98%. When using several mirrors like in Figure 16 on page 27, the cumulative loss of brightness might become noticeable.

Glass mirrors also have the advantage of having better rigidity, but that is usually not an issue with a CAVE system, where the frame can be built rigid enough to hold the mirror in a single plane. In some odd configurations, this might also be an issue, though. One example case is that loud sounds or air currents may induce vibrations to the foil mirrors, disrupting image. This can be alleviated somewhat by providing a back plate that keeps the foil mirror more rigid.

Mirror dimensions

When calculating the actual mirror width required one must remember that a CRT projector consists of actually three different optical systems, and that their optical axes do not originate in the same place, though they do converge at the screen (assuming the projector is adjusted correctly). See Figure 27.




Figure 27: CRT projector light cones and convergence.

The true image width on the mirror may be easily calculated using basic geometry. The formulas are omitted here, please see *Appendix A: Projector Geometry* on page 107 for a detailed discussion on the image geometry.

Note that in general it is useful to reserve some extra space at the mirror edges, because the edges may not be entirely planar. Especially with foil mirrors the image should not be projected closer than 5 cm to the mirror edges. Also, the mirror frames may take a couple of centimeter off the available mirror surface area, but naturally this is not cumulative with the previous comment.

2.3 Projectors

2.3.1 Display Resolution

The physical resolution of the projector is very important. Today you often see presentations where the Windows desktop is projected from a laptop to a screen, and the result is completely unreadable, because the true physical resolution of the projector is 800x600 pixels, but it is trying to display 1024x768 pixels by omitting pixels at certain intervals. This is called downsampling. The image quality may be enhanced by performing downsampling with anti-aliasing, but regardless of the way it is done, some image quality is lost.

When we wish to display a 1024x1024 image on the projector for VR purposes, we need to take notion of the fact that the CAVE walls are square, but the display area of the projector is not, that is, their aspect ratio is not the same. See Figure 28.




Figure 28: Losing physical pixels due to aspect ratio.

As can be seen from Figure 28, the square screen loses 1/4 of the pixels available in the vertical resolution, which means that in order to display a 1024x1024 image the minimum physical resolution is 1365x1024 pixels.

In practice, however, in CRT projectors the image should be kept away from the edges of the phosphorus to avoid burning [Barco 97a], which in turn loses approximately 15-20% of the available pixels horizontally and vertically, turning the required resolution to 1600x1200 pixels. Of course, more is preferred. The best currently commercially available CRT projectors have 2500x2000 pixels true resolution [Barco 98], [AmPro 97]. With LCD projectors there is no reason to keep the image away from display area edges, so less true resolution is required.

2.3.2 Display Refresh Rate

Perception of Flicker

The eye integrates the light intensity over a short period of time, which can extend for several tenths of a second. Since in video display devices images are displayed sequentially, the eye is able to see the refresh, if it is too slow. The limit where the eye can no longer see individual images, known as Critical Fusion Frequency (CFF) can be approximated using the well-known Ferry-Porter Law [Kalawsky 93]:

$$f_c = a \log B_a + b, \quad (2-2)$$

where a and b are adaptation constants ($a = 12.5, b \approx 37$ for the photopic¹³ case, and $a = 1.5, b \approx 37$ for the scotopic¹⁴ case), and B_a = luminance of test field. Unfortunately, most forms of display devices do not turn their image off immediately, which makes determining the exact CFF difficult. The CFF is also dependent on the displayed area ($f_c \sim \log A$). Generally speaking, refresh rates above 50 Hz are good approximations for larger screens, such as a CAVE. Typically used values are 50 and 60 Hz, corresponding to the PAL and NTSC field rates. Movie screens have a frame rate of 24 Hz, but a frame is shown twice to reduce the flicker, making the effective display refresh rate 48 Hz. Being mechanical in nature, they also have an almost zero image turn-off time.

Displaying Stereoscopic Images

One of the key characteristics of a stereo-capable display is the rapid refresh rate, which varies from 100 to 120 Hz (50 - 60 Hz for either eye). It is not enough that the projector is able to sync to this refresh rate, but it must also be fast enough to remove the traces of the previous image before a new one is shown, that is, the phosphor persistence time may not exceed the duration of a single image. If this were not the case, and the stereo effect was achieved by displaying the left and right eye images sequentially, the user would see “leakage”, or “crosstalk” from the other eye, which would cause a partial loss of depth effect.

Slow persistence time would also mean that fast-moving objects would be subjected to ghosting and other artifacts. This includes of course full scene transformations, such as when the user moves about in the virtual world.

Image brightness may also be a contributing factor: with brighter images the phosphor takes a longer time to “cool down”, resulting in crosstalk [Kalawsky 93].

13. Bright surroundings, color vision available. This is the main state inside a CAVE.

14. Dark vision.

In a 3-tube CRT projector the green phosphor is usually the slowest, which may cause leakage. This is why the projector should be equipped with a “fast phosphor”-option, which unfortunately tends to dim the green CRT slightly, causing an overall drop of intensity.

LCD, as well as lightvalve and micromirror device projectors are not usually able to change their internal state more than 60-80 times per second, which make them effectively unusable for field sequential stereo. Note that this ability has no relation to which vertical refresh rate they actually are able to synchronize to: many LCD projectors claim to sync to 100 Hz or more, when in fact their internal refresh rate is usually lower and some sort of downsampling takes place. LCD projectors are, however, quite usable for field interlaced stereo or polarization-based stereo, since these techniques do not require faster refresh rates than standard monitors.

2.3.3 Bandwidth

Stereo projection requires that the image needs to be refreshed at very high rates. In addition, we are using high resolutions for the image (approximately 1000 x 1000 pixels). All this means that we are using a lot of image *bandwidth*.

The approximate RGB bandwidth may be calculated as follows [Extron]:

$$B = 1.5 \times R_x R_y f \quad (2-3)$$

where B = bandwidth, f = field refresh rate and R_x and R_y are the horizontal and vertical field resolutions, respectively. A more accurate calculation would take into account the different characteristics of a video signal, such as vertical sync time, back porch, front porch, and so on, but the approximate formula of Equation 2-3 is quite sufficient as a rule of thumb. A detailed discussion is available in [Eitzmann 98].

For a typical 1024x768 display at 120Hz we then get the RGB bandwidth to be approximately 141 MHz, which is at the very top of the most high-end projectors. For example, the quoted RGB bandwidth for the Barco 1208s data projector is 120 MHz [Barco 98].

If the data rate exceeds the RGB bandwidth of the display device, the projector will start losing resolution and the image will begin to smear. Thus the refresh rate or the resolution should be kept below the minimum allowed by bandwidth. In practice, this limit may have to be somewhat exceeded since with large displays a high resolution must be used. However, the image quality reduction is acceptable, unless very high detail is used. Also, pixel resolutions above 1024x1024 are practically useless for stereo projection, because the required RGB bandwidth (180 MHz at 114 Hz) is more than current display technology is able to handle.

It should be noted that the bandwidth issue does not only apply to the projector, but the intermediate cabling as well. As a rule of thumb, you should use only equipment with at least twice the bandwidth you are going to use, since each piece of electronics and cabling will cause a loss of signal [Extron].

2.3.4 Geometrical and Other Distortions

As already has been shown, there are multiple possible configurations for a CAVE system. Not all of these configurations have the projector in exactly 90 degree angle to the screen, nor may the projectors optical axis go through the exact screen center. The screen may also be slightly curved, due to construction and stretching, nor may it be entirely right-angled.

All these defects mean that the projector must be able to correct for a number of different geometrical and other distortions [Vance 87], [Barco 97a]:

- **Non-uniform focus.** When the projection plane is tilted, the distance to different parts of the screen varies, causing some areas go out of focus.
- **Keystone distortion.** Areas closer to the projector are smaller than corresponding areas that are farther away. This variation in horizontal magnification causes keystone distortion.
- **Vertical and horizontal non-linearity.** The variation in the vertical magnification also causes vertical non-linearity. Similarly for horizontal non-linearity.
- **Bow, skew, and seagull.** These are caused by non-linearities of the screen, and show as curved or tilted image.
- **Convergence.** This is an artifact of multi-tubed CRT projectors. See *Cathode Ray Tube* on page 48.
- **Color adjustments.** Typical color adjustments that are required are the color temperature, brightness, contrast, and gamma. These are important when matching the colors of the different CAVE screens.

Other, less used distortions also exist, such as the Scheimpflug (diagonal focusing, [Barco 97b]), but they are skipped here for brevity.

The selected projector should provide a large array of controls to compensate for these distortions, preferably with an automatic setup to reduce the adjustment time. Manual control will be required often, regardless of any automated system available, so controls should be easy to use and logical. Commentary from different CAVE administrators around the world indicate a mean time from a week to a month before the projectors have to be readjusted, with full calibration done every six months. The time depends on the required image quality and the chosen construction method. There seems to be no real difference between projector providers.

2.3.5 Projector Types

Large screen projection has advanced rapidly in the past few years, mainly thanks to the advances in integration technology. The old Cathode Ray Tube (CRT) projectors are being largely superseded in different applications, but CRT technology is in general still usable for multiple purposes.

The following is a short write-up and comparison of different projection technologies, based mostly on [Extron 99a].

Cathode Ray Tube

Cathode Ray Tube (CRT) displays were made popular in household TV sets where they are still commonly used, but have been declining in popularity for the past few years especially for computer graphics applications.

The CRT technology is old, and as such, has many problems:

- CRT projectors use commonly three optical systems to achieve color (RGB). This means that in order for a projector to be calibrated, the different optical systems must *converge* at the same point, which makes both installation and maintenance complex. Other geometrical corrections become more difficult with three optical systems instead of one.
- The three separate color guns require careful adjustment to make sure the color balance is correct. With multiple projectors, all must be adjusted correctly.
- The CRT phosphor has a limited life. If multiple projectors are used, care must be taken that they all are used the same amount to minimize differences in image quality.
- CRT projectors require space, because of the three separate optical systems.
- Projectors are expensive (compared to other types), since three tubes and lenses are not cheap to build.
- CRT projectors are not very bright (in the order of 200-300 ANSI lumens). Note however that brightness is not an entirely desirable feature, see 2.2.3 *Brightness and Contrast* on page 32.
- The contrast ratio is approximately 1:100, meaning that the brightest white pixel they can project is 100 times brighter than the “black” pixel.
- While high bandwidth CRTs are available, the top models achieve only 120 MHz. On the other hand, being analog, they degrade gracefully when the image bandwidth exceeds the projector bandwidth.

On the other hand, CRT technology is well known, and projectors are readily available. In addition, the CRT projectors are able to change their state very rapidly, and can sync to very high refresh rates, which allows field sequential stereographic projections to be used.

LCD

Liquid Crystal Display projectors are based on the standard LCD principle: the light is passed through LCD panels that can control how much light they pass. LCD projectors are:

- + Light and cheaper to manufacture
- + Are in general brighter and have better contrast ratio.
- + Have only a single lens, which makes user convergence adjustments unnecessary.

However:

- The LCD projectors do not have very good image quality, nor do they have the capability for very high resolutions (at the moment).
- LCD crystals are slow, and thus at high refresh rates images smear, badly, making them unusable for field sequential stereo displays.
- Color is achieved by placing color elements close to each other, known as triads. With big screen size, they may become visible at close ranges, diminishing the image quality.

DLP

The most common DLP or “Digital Light Processing” systems are Digital Micromirror Devices (DMD), which are based on small mirrors that change their state. They have numerous advantages over standard systems and are superseding LCD projectors:

- + DLP systems can be made very light and cheap.
- + The same optical system can be used to project all primary colors, making convergence adjustments unnecessary.
- + Digital technology evolves fast, and DLP projectors become faster and better with each generation.
- + Brighter images than with standard CRT projectors, up to 1000 ANSI lumens.
- + Better contrast ratio: up to 1:400 are available.
- + Higher bandwidth: some DLP projectors have a 150 MHz bandwidth.

However, there are serious drawbacks that make using them in a VE projection system difficult:

- Low resolution. Even the better DLP systems cannot do very high real resolutions. The high advertised resolutions are done by downscaling the incoming image, reducing image quality.
- DLP systems cannot change their internal state very rapidly, limiting the effective refresh rates to 60 Hz.

- The lamps used with the DLP projectors need replacing often, from 40 hours to 1000 hours, depending on the technology.
- The produced picture is discrete, as opposed to the continuous picture of a CRT. This prevents the gradual degradation of the signal at high bandwidths and the image elements may be seen individually on large screens at close distances.

Light Valve

Light valves can be built with multiple technologies, either using LCD displays or DLP processors. They produce very bright images, but have the same advantages and disadvantages than the LCD or DLP systems, respectively. The contrast ratio is usually better than 1:100.

One distinct version of Liquid Crystal Light valves is called Image Light Amplifier, which was first developed by Hughes aircraft systems in the 1970s [Hughes 99]. Three small CRTs are used for initial image generation, and the light is amplified using a liquid crystal layer and a powerful lamp. This results into a very bright (up to 12000 ANSI lumens) image, with all the benefits of a CRT, except for speed: a typical ILA can cope with refresh rates of up to 80 Hz, but no more, the LCD technology being the limiting factor. The contrast ratio is also excellent, up to 1:1000. However, the setup process and convergence adjustments are more difficult with an ILA projector than a CRT projector. The ILA systems are also rather expensive.

2.4 Tracking Systems

Tracking applications are currently developing rapidly, as the entertainment industry has found out that motion capture to be an efficient technique for computer animation, resulting in big demand for fast, accurate, and cheap motion tracking devices.

Tracking is usually done either actively or passively, though combinations of them do exist [State 96], [Azuma 98], [Auer 99].

2.4.1 Tracking Systems Properties

Accuracy

Accuracy of the tracking system can be divided into static and dynamic accuracy. Static accuracy determines how accurately the system can give the location of a static object, whereas dynamic accuracy reflects how well any changes in the objects position or orientation are measured. In a CAVE environment the static accuracy should be as high as possible.

Update rate

A major contributor in the dynamic accuracy is the update rate, or the bandwidth, of the tracker. If it can produce accurate data once in a minute, it is clearly not suitable for a highly interactive task, such as movement in a Virtual Environment.

Latency

The problem with tracking isn't usually the update rate, since the sensor and the computer often are connected with fast local area networks (LAN). The problem is latency, which means the time which it takes for a packet of position information to be acquired and be processed before it is applied against the image. Big latency can in extreme cases induce motion sickness, if the projected image lags significantly from the actual user position [Burdea 96]. The effect is more pronounced with HMD devices.

Latency can be combatted using several means, by simple hardware configurations (like using direct serial or optical networking between peers) to more complex, predictive software filters [Wu 95].

Registration

Registration is the correspondence between the actual position and reported position. This is not the same as accuracy, since poor registration may be caused by accumulated errors. In a fully virtual environment this is easier to ignore than in an augmented reality application, where any discrepancy between the VE and real world is easily visible.

Noise

A good tracker should have a good signal-to-noise ratio (S/N). A tracker that "jitters" badly has to be compensated for in software (by filtering) which in turn increases latency and phase lag. Predictive filters may then be used to reduce the latency [Wu 95].

2.4.2 Tracking Hardware

There have been many surveys done on the available tracking hardware [Ferrin 91], [Kalawsky 93], and so this presentation is only cursory, with evaluation on how they could actually be used in a CAVE.

Mechanic Tracking

Mechanic tracking is probably the fastest and most accurate of the tracking systems presented here. Unfortunately, the required hardware is not very useful in a CAVE environment, where free movement is desired. For a mechanical tracker system, see the FakeSpace BOOM in Figure 4 on page 14.

Magnetic Tracking

Tracking systems based on detecting changes in magnetic fields are currently probably the most used systems. Companies such as Ascension and Polhemus have introduced medium-cost, very accurate tracking systems that are easy to use and not very difficult to wear. Kalawsky gives an excellent account on the magnetic trackers in [Kalawsky 93].

Magnetic tracking may be either AC trackers or DC trackers. AC trackers use a changing magnetic field, and thus induce eddy currents into any conducting materials nearby, which then cause secondary magnetic fields, disturbing the measurement.

DC trackers on the other hand use a pulsed DC magnetic field, which allows the eddy currents and therefore the secondary magnetic fields to settle down before measurement. This is very useful when there is a large amount of conductive metals present, as they do not impact the performance of the system (as much - there are of course residual currents, especially with ferromagnetic material). The calculations do require more power and more complex hardware, though, and the system is still vulnerable to ferromagnetic metals, since they generate their own magnetic field. To counter this the Ascension DC tracker measures also the ambient magnetic field for self-calibration.

Inertial Tracking

Inertial trackers are very good and useful in the sense that they are sourceless and operate over an unlimited distance. The trackers are also cheap and easy to manufacture in large numbers, and they are readily available from manufacturers such as VTI Hamlin in a variety of packages [VTI Hamlin 98].

The main problems with accelerometers is the drift inherent to the required integration. To get speed from acceleration you have to integrate once, and to get location, twice. This integration accumulates measurement and round-off errors very rapidly, and the calculated location drifts away in a few seconds.

Some of the errors may be countered by using the Earth's gravitational field as a reference, since its magnitude and direction is well known. However, this is only useful when tracking orientation, and even then it can only counter two of the three axis (yaw cannot be deduced from gravity). The relative cheapness of the sensors does make them attractive options for applications where absolute location tracking is not necessary, but can be replaced with relative location [Ilmonen 00].

The drifting problem is often solved by using a secondary tracker system which calibrates the data received from the inertial sensors. Typical systems are optical, RF, or ultrasound [Foxlin 98].

Commercial inertial tracking systems are currently available: InterSense, Inc., has developed Constellation, an inertial tracking system, which uses low-cost accelerometers to track user's movements, combined with ultrasound system that is used to calibrate the system on-the-fly [Foxlin 98].

Acoustic Tracking

Acoustic tracking can be done in two ways: measuring the Time-of-Flight (TOF) between an emitter and a receiver, or measuring the phase difference (phase-coherent, PC) between the emitter and the receiver. Ultrasound frequencies (over 20 kHz, well outside human hearing) are used.

The data rates for a PC tracker are better than for a TOF device, because phase can be constantly measured, but the time-of-flight in a typical room takes several milliseconds anyway, making latency a big issue for acoustical trackers. In addition, acoustical trackers tend to suffer from air currents and occlusion.

By using multiple sensors attached to the tracked object, and using measurements such as the time or phase difference between sensors, it is possible to measure the orientation as well [Foxlin 98].

Optical Tracking

Optical tracker systems are much harder to make than others, since there usually is very much redundant or extra information in the scene. Also, physical constraints such as occlusion by other objects is a problem. The object may also be difficult to discern in certain positions or lighting conditions, making computer vision applications hard [Sonka 93].

The rewards are great, though: no wires, no sensors, no extra equipment; automatic gesture and full-body posture recognition; ease of use. It is no surprise that a lot of work has already been put into this field, and with the computing power now available to us, a lot of results are emerging [Dorfmueller 99], [Park 99].

Typical methods through which optical tracking is achieved is through using different markers [Starner 98], [Kato 99], [Pietarila 00], using a similarity measure with pre-processed images [Aoki 99], pose recognition and computation [Park 99], etc. Common to all these methods and optical tracking in general is that they require plenty of computing power, which is the reason passive optical tracking hasn't been pursued actively until now. Active tracking has been done longer, where either the user is wearing LEDs or IR transmitters, or a camera which tracks LEDs or IR transmitters in the room [Ward 92].

For a CAVE-style application optical tracking is somewhat more difficult: usually the person working inside the CAVE is not alone, which makes the task of the optical tracker more difficult, when people move around inside the CAVE and cause occlusion. Also, the physical structure of the CAVE (the walls will be in the way) and constantly changing light conditions cause problems.

Other Tracking Methods

Some innovative tracking methods have been developed recently. For example, one such system determines the user's location in a building using a combination of gravity/temperature/air pressure sensors [Golding 99], though not very accurately. Also, GPS locators and digital compasses have been used outdoors for determining the location of the user [Piekarski 99]. However, the inherent inaccuracy in GPS makes it unsuitable for precision tracking applications even though the accuracy can be enhanced significantly by using differential GPS. The indoor nature of CAVE environments makes GPS unusable, because of the required clear visibility to the satellites.

2.5 Computing Equipment

2.5.1 Main Graphics Computer

The typical CAVE graphics computer is nowadays built around the Onyx2 architecture from SGI (former Silicon Graphics) [SGI 98]. This high-end model can be equipped with multiple graphics options, allows a high-end scalable system ccNUMA architecture enabling up to 128 CPUs with almost linear increase in processing power.

The graphics subsystem on the Onyx2 (usually InfiniteReality or InfiniteReality 2, which is the newer and faster version) is very versatile, and is able to output its power to a very fast single display, or split it across several channels. Up to 16 graphics subsystems (referred to as "pipes") can be installed into an Onyx2 computer. For a more detailed discussion, see *Components of the Rendering Pipeline* on page 55.

The Onyx2 combined with the InfiniteReality graphics has a number of additional, useful features:

- Support for large drives (2.3 TB internal storage) and large memory (up to 256 GB).
- High quality graphics output, including full scene antialiasing (see below for more explanation) at very high polygon rates (13.1 million triangles/sec. for a single pipeline), fast texture download, clipmapping, projective textures, shadow maps, video to texture, atmospheric effects, dynamic resolution, and so on. For a full list, see [SGI 98]
- High graphics bandwidth: 5.6 GB/s for fastest combination.
- X window system, which can be distributed over multiple displays
- SGI's solid experience on high-end graphics
- Compatibility: many CAVEs around the world use the same hardware, making applications easily available.

Let us now review some of the more important components of the Onyx2 InfiniteReality system.

Components of the Rendering Pipeline

The current top subsystem is InfiniteReality2, or IR2¹⁵. The IR2 graphics pipeline consists of three main components: the Geometry EngineTM (GE), the Raster Manager (RM), and the Display Generator (DG).

The Geometry Engine processors do geometry calculations (such as transformations, lighting, clipping, and projection to screen space) and pixel processing (including different image processing operations) for pixels, textures, or video.

The Raster Managers convert the triangle, line, and point data into pixels applying such effects as anti-aliasing (including multisampling), and texture mapping. RMs also provide raster memory (64 MB/RM). Between two to four RMs can be installed into a pipeline.

The Display Generator is responsible for converting the digital framebuffer data into analog signal that is sent to the display device. The DG5-2 can drive two channels, whose resolution depends on the number of the Raster Managers available, whereas the DG5-8 can drive up to eight channels.

Since the DG system is able to provide multiple channels, it is possible in theory to build a full CAVE system using only the minimum possible configuration with one pipeline. However, since the geometry processing power is split up between the different channels, and the RM frame buffer memory is limited, multiple pipelines should be used. A typical four-wall CAVE configuration has two pipelines that are split in half, for a total of four channels. Usually each pipe has at least two RMs, because typical VR applications need more texture processing power than geometrical processing power. Geometrical processing power can be added by purchasing more GEs, which in turn requires more pipelines. The DAC on the DG is slow (220 MHz), and thus resolution of multiple images from a single pipeline is very low.

The DPLEX (Digital Multiplexing) mode allows multiple IR pipelines to drive a single channel. A 16 pipeline RealityMonster can thus drive a four-wall CAVE so that each wall receives the full graphics capability of four pipelines.

Multiprocessing

While many basic graphics operations are done by the display hardware, CPU power is still required for application usage. The Onyx2 architecture allows up to 128 MIPS R10000 or R120000 CPUs to be installed, while the architecture has been designed to allow for almost linear scalability.

15. New architecture is expected from SGI in H2, 2000.

CPU power is not only required for application use: some graphics tasks can also be assigned to the CPUs, such as culling (selecting which objects to render and which to discard), which can be very computationally intensive. Drawing is usually done in a separate thread, while the scene graph is computed in the application thread. This is done to ensure that the display update can be kept as constant as possible.

A VR application thus can use at least three threads (four with stereo) for each graphics pipeline: Application, Cull (one for each eye), and Draw, and thus the recommended number of CPUs per pipeline is four. In practice, however, the application thread is often shared among pipelines and the Cull and Draw threads often do not use all processing power, and so two CPUs per pipeline is often acceptable. Using more than four CPUs per pipeline does not add performance, unless the application can be parallelized effectively.

Resolution & Antialiasing

A typical CAVE screen is 2.5 – 3 meters in width and height. In a typical application, the resolution is 1024x768 pixels, which translates to 3-4 mm/pixel on the display medium, which makes the pixels very clearly visible, resulting into jagged edges. The analogue nature of the projector usually smears the pixels somewhat, as also does the fact that proper focusing is hard to do on a back-projection screen. However, since the available bandwidth (see *Bandwidth*, page 46) limits the usable resolution to no more than 1024 pixels square, the only viable choice is to antialias the imagery [Foley 90].

Since antialiasing is rather expensive to do in software, current high-end graphics hardware offer a no-cost solution [Akeley 93], [Montrym 97]. With the 8 pixel random subsampling from an 8x8 grid in the InfiniteReality architecture from SGI, the computational resolution of the display grows by 180%, transforming a 1280x1024 display into a 3620*2896 display. Of course, only 1280x1024 pixels are shown, but sampling is performed in a 3600x2900 pixel grid, since choosing eight random subsamples in an 8x8 grid averages 2.8 pixels horizontally and vertically.

Operating System and Support

All nice hardware is wasted, unless the operating system can fully utilize the multiple CPUs and graphics pipelines. The 64-bit IRIX 6.5 Operating System that comes with the Onyx2 is fully scalable SMP (Symmetric Multi-Processing) OS.

The IRIX scheduler supports real-time processes and batch processes, in addition to normal user processes. Real-time processing capability combined with low system overhead is useful in real-time graphics: A guaranteed maximum latency allows the application to respond quickly to events, and make services like guaranteed frame rate [Eckel 98].

IRIX also supports shared memory, allowing fast communication between processes, even those running on different CPUs. This allows for example easy scene graph double buffering.

SGI also offers both low-end (OpenGL [Neider 93]) and high-end graphics libraries (IRIS Performer [Eckel 98]) that have been tuned to maximum performance on their systems. Both libraries are currently in wide use; OpenGL as a generic cross-platform 3D library and Performer as the high-end visualization library.

One important thing that should not be overlooked is that at the current rate of evolution of computer technology, even high-end systems are overrun every few years. However, it is economically not feasible to keep changing the hardware, because at high-end it tends to be very expensive. It is thus important that the company is able to give support to the software and hardware for several years, and possibly a graceful upgrade path.

PC-based CAVEs

There are not very many PC-based CAVE-like systems available. At the writing of this only one commercial solution is available, from MechDyne corporation.¹⁶

Building a VR system using standard PCs does have its advantages: the hardware is cheap and price/performance ratio is very good. Multiprocessing solutions which are a must for any serious VR facility are also available, though at higher costs than off-the-shelf Windows boxes. There is also a great variety of software available, and a good deal of it is cheaper than the corresponding UNIX equivalents.

However, there are multiple negative points which make building PC-based CAVEs quite difficult.

- Multi-processing support is currently very scarce and expensive, when available. Most hardware suppliers support at most 4 CPUs and while larger arrays do exist [Warren 98], [Chien 99], the operating systems often do not have the advanced capabilities required to fully utilize such network, and as such, these are better suited for numerical processing than real-time computer graphics.
- Operating systems are not in general tuned for maximum graphics performance, and often do not have the capabilities to handle real-time processes, nor efficient shared memory architectures. (Games often bypass the operating system somewhat by using their own graphics engines)
- No multiple graphics pipes are available. The ability to plug in two graphics cards is not equivalent, since the pipes must be synchronized. In addition, two cards is not enough to drive a four-wall CAVE. Using things such as video splitters or multiple computers this is possible to circumvent, though,

16. <http://www.mechdyne.com/>

but at the cost of image quality or increased complexity.

- There are no established high-end graphics standards such as the IRIS Performer in the PC/Windows¹⁷ world. While this does not prevent VR development, it does make the life of a developer more difficult. In addition, there is currently a standards war between OpenGL and Direct3D as a lower-level standard.
- The ever-changing PC hardware world, while simultaneously a blessing, is also a very difficult area to support: manufacturers are often more interested in making faster hardware than providing correct driver support. This makes it difficult to support even a reasonable subset of available hardware, so incompatibilities may occur frequently.
- Some PC manufacturers (such as InterGraph and the SGI with their late Visual Workstations) offer their own, specialized 3D hardware. While this allows a much better driver support, it also means that the manufacturers have to constantly battle against specialized companies that do nothing except design new chips, meaning lower margins and lower profits.

The standard PC has found its use in desktop virtual reality where cheap power is often more desirable than rendering quality or multi-processing support, and they make fine auxiliary computers to augment a CAVE facility.

Interestingly enough, the Apple Macintosh has been used in providing spatially immersive displays for some time now [Panoram]. Though not usually regarded as a visualization computer, the operating system and hardware support for multiple displays has made it a useful low-end choice for panoramic screens. Lack of support for suitable software and hardware precludes it from being used for stereoscopic projection, though.

2.5.2 Auxiliary Computers

Auxiliary computers are needed as pre-processing computers, handling audio, and development and testing of new software. If possible, they should thus support the same APIs and preferably operating system as the main computer. For the SGI Onyx2 the auxiliary computers usually are of the O2 and Octane variety, depending on whether they are used as personal workstations or R&D computers.

17. IRIS Performer is available for Linux since 29-Nov-1999. See <http://www.sgi.com/> for details.

2.6 CAVE Software

This chapter will look at the basic requirements CAVE software should be capable of, and then a short look at the currently available software is provided. Its software selection is limited to freely available, CAVE-specific software. High-end graphics libraries such as IRIS Performer and Open Inventor are omitted, because they are not CAVE-specific, and a substantial amount of coding is required to make applications run on the high-end libraries. A more comprehensive look at available software is given in [Bierbaum 98] and [Laakso 99].

2.6.1 Basic Requirements for Software

There are a few basic requirements that a CAVE software package should be capable of:

- Ability to load, display, and synchronize a 3D scene in a multi-walled display environment.
- Ability to plug in different controller types such as a wand or a mouse.
- Ability to provide hooks or other methods by which objects can be manipulated, changed, added, or removed.
- Ability to display a stereo image using the methods described in Chapter 1.3: *How Three-Dimensional Displays Are Made?* on page 8.
- The software must be configurable to different hardware configurations easily.

Cavelib

The Cavelib is the original EVL Cave library that is now distributed by FakeSpace/Pyramid Systems [Cruz-Neira 95]. It is built on both OpenGL and Performer in C, and provides rudimentary drivers for the CAVE display and other hardware.

The Cavelib is not very advanced, but it does have the edge of being the most compatible.

AVANGO

AVANGO is a new VR support library available from GMD, Germany [Dai 97]. It has been designed from ground up to be fast and flexible, and has been built on a top of curious combination of C++ and Scheme (a Lisp-variant, [Abelson 85]). AVANGO was formerly known as Avocado, but the name was changed due to trademark problems.

AVANGO duplicates the IRIS Performer scene graph, and supplies a number of its own nodes for audio, and other things. This makes it very tightly dependent on the Performer libraries, and while theoretically portable, it is unlikely to be available for other platforms (with the exception of Linux). It even supports olfactory interfaces [Göbel 00].

AVANGO is not freely available, but test licences can be requested from GMD. The future of AVANGO is not known at the present time.

VR Juggler

The VR Juggler is a new development by the Electronic Visualization Lab, where the original CAVE system and the Cavelib was developed [Bierbaum 98], [Just 98]. VR Juggler is designed to overcome the previous Cave libraries by establishing a common, modular architecture for different devices and spanning the whole VR scheme using a single system.

VR Juggler assigns the abstraction of different hardware (trackers, input and display devices, windowing systems) to entities called Managers, which are governed by a single Kernel. The application then talks to these managers and never learns the truth about the underlying hardware, which may disappear, be changed on the fly, or be dynamically reconfigured. The only layer the application talks directly to is the graphics API to maximize performance.

VR Juggler is written in C++ and is freely available from <http://www.vrjuggler.org/> for SGI IRIX, HP-UX, Linux, and Windows NT. Currently supported graphics APIs are OpenGL and IRIS Performer. The software is being actively developed and slowly moving out of the academic world.

Maverik

MAVERIK (Manchester Virtual Environment Interface Kernel)¹⁸ is a C toolkit for interactive VR applications. It is a very recent development, but unfortunately it is also in a very infant status. The Maverik is cross-platform, C and OpenGL based library, which does hold some promise for the future, but at the moment the render quality and the hardware support (no real stereo, for example) is completely inadequate for anything else than running simple VR environments in a HMD system. Maverik is described better in [Laakso 99].

18. <http://aig.cs.man.ac.uk/systems/Maverik/index.html>

2.7 Audio Systems

Sound has often been neglected in VR systems. Most installations provide only stereophonic sound via a headset or surround sound via 4- or 8-channel loudspeaker arrays, usually set in the corners of the CAVE.

2.7.1 Acoustics

As darkness and absorbing, dark material is required to keep ambient light from disturbing the virtual image, so does the sound need to be controlled for the synthetic VE sound. The room's own acoustics need to be cancelled and the room must be made as anechoic as possible [Lokki 99], because room equalization cannot be performed for multiple listeners, especially when they are moving.

However, building a fully anechoic chamber is often impossible: A hard floor is needed for walking, the projectors may provide a source of noise by being in the same room, and air conditioning may provide additional noise. As much noisy equipment (such as the computers) as possible should be placed in a separate room.

The noise coming from the outside world should be eliminated as well, as should the sound from within the CAVE to the outside world. The latter is often forgotten, but is increasingly necessary as noise becomes more of an environment problem in the modern world.

The effect of the equipment inside the room is difficult to quantify. While the smaller objects such as projectors are not very effective diffusers, the large areas such as the mirrors (which, being hard, tend to reflect sound) and the screens (soft screens absorb the sound, hard screens reflect) affect the global acoustical environment in a very frequency-dependent way. No research is known that addresses this matter.

2.7.2 Sound Reproduction

Binaural - HRTF

Binaural sound reproduction is usually made using earphones, so that the sound arriving at both ears can be controlled with no fear of crosstalk. The prevalent method of generating binaural sound is using HRTFs (Head-Related Transfer Functions). A good introduction to HRTFs can be found in [Huopaniemi 99].

HRTF design and implementation is demanding. In addition, for each user of the virtual room an own set of headphones must be provided, and calculations should be done for each user individually to provide a best possible experience.

Binaural loudspeaker sound reproduction is not recommended, because 3D sound reproduction cannot be done completely with only two loudspeakers.

Multichannel - Ambisonics

Ambisonics is a spatial audio encoding technique where the 3D sound field is encoded into four channels. However, this technology is best suited for reproducing recorded sound fields, and does not suit well for dynamically generated virtual environments, since it is very sensitive to the listeners location.

Multichannel - Vector Base Amplitude Panning

Vector Base Amplitude Panning (VBAP, [Pulkki 97]) allows the arbitrary positioning of loudspeakers. The loudspeakers are placed as a triangle mesh over a sphere (optimally), and the sound is panned using the VBAP technique, reproducing a fully 3D sound. VBAP installations should use at least 8-16 loudspeakers, but for best effect, more loudspeakers may be used.

VBAP can be easily integrated with different acoustics models to provide an acoustic simulation of the VE as well [Savioja et al 99], [Savioja 99]. Since the sound is audible from three loudspeakers at any time, there will be a small error as the listener moves around. However, it is not very perceivable, since in typical CAVE installations the user movements are limited and the loudspeakers are placed outside the screen array.

Loudspeakers are not without their problems: the positioning and alignment of the loudspeakers is critical, and the screens provide dampening, which is frequency-dependent. Also other objects which cannot be made anechoic such as the mirrors, the projectors, and possibly the floor, will cause reflections that may be impossible to compensate for, especially if the user is moving around over a large area (compared to the positioning of the loudspeakers).

2.7.3 Other Sound Sources

EMFi and other loudspeakers

The normal loudspeakers could possibly be replaced with electromechanical film (EMFi), which would allow the construction of much larger loudspeaker arrays, and make the spatialization of sound easier [Antila 99]. It remains to be seen whether this new material is more suitable to immersive displays than standard loudspeakers.

Vibration Emitters

To provide the users with very low frequency stimuli, some CAVE systems, such as the CyberStage, have installed vibration emitters below the floor [Eckel 99]. They are used to provide a more immersive feeling, for example collision shocks, etc.

2.8 Infrastructure

2.8.1 Cabling

Layout

A fully loaded CAVE with six projectors, a score of IR transmitters, 16 loudspeakers and an eight-channel tracker system generates a lot of wiring (for the stated case, 23 power cables and 50 data cables, at least). Care must thus be taken on how the cabling is handled.

Preferably all cabling should be non-intrusive and tucked someplace safe, yet easily accessible for maintenance and replacement. The cable layout should be logical, and naturally, all cables should be identified clearly at both ends, possibly even in the middle. Typical solutions include a central rack, through which all necessary data cabling is routed. The cabling is then routed via cable corridors on the wall or in ceiling to the equipment. Floor placement should be avoided because it is too easy to step on the wires. If floor placement is necessary, the cabling should be drawn through one route only and protected with a cable corridor.

Video and audio cabling

The graphics computer outputs a lot of data for a CAVE environment. In a display chain, the weakest link defines the maximum throughput of the entire chain, and thus the cabling and any switches or other equipment should be of high quality. Typical bandwidths that are passed through the video channel are in the order of 140 MHz (see 2.3.3 *Bandwidth* on page 46). A rule of thumb says that the bandwidth of intermediate equipment should be roughly twice the required bandwidth.

If cabling becomes very long, signal attenuation needs to be taken into account. It is commonly expressed as dB/m or dB/100m. For very long cabling, one should consider optical fibre, since it has very low attenuation, and complete immunity against RF/EM interference [Extron 99b]. Optical video/keyboard/mouse transmitters are available for example from Lightwave Corporation [Lightwave 98]. For maximum cable length consult the manufacturer of a particular cable/device.

Power

There are two important things about power:

1. There must be *enough of power*. Brown-outs are very dangerous to electronics, especially expensive electronics. Note that in case of a black-out, all equipment will be starting up at the same time when power is available again, and typically the peak power consumption is significantly more than average consumption.

2. All equipment electronically connected must *share the ground*. If the ground level between, say, the monitor and the computer is different, there will be current in the ground wire, which in the worst case may break the equipment.

Other than that, the power cabling should be treated with the same respect as data cabling and be drawn the same way. Not necessarily through the same corridors, though, since the AC current may in some cases cause interference in small-voltage DC cabling. If a magnetic tracker is used, cabling should not be drawn near the emitter unit, since the copper and AC current will cause interference.

2.8.2 Layout

In general, it is useful that the CAVE operators occupying the console have visibility to the inside of the CAVE. This makes communication easier, and allows the CAVE operator to dynamically change the environment, if need arises. It helps if the CAVE screens can be duplicated on the operator's monitors, either through a software solution, or then by splitting the display signal. The IR pipeline allows a standard PAL/NTSC signal to be generated from the current frame buffer, and then be shown through cheap TV monitors.

Light is also required in the room to some extent, which clashes with the need for total darkness. The required light for the operators can be provided by tabletop lights, or better, the ambient light level on the room should be controllable to allow for low enough light levels for reading, but not too high that would diminish image quality. The lighting at the operators end of the room can be made separately adjustable, and the CAVE display systems separated using a wall or a curtain.

Separation of the operators space and the equipment space is recommended also for the purpose of keeping unwanted visitors from disturbing the delicate mirrors or projectors. Even if direct contact is avoided, dust and dirt are easier to keep away when users have no access to the equipment area. Care should be taken that separating the two areas does not disturb the acoustical qualities of the room. Local fire regulations should also be observed.

2.8.3 Miscellaneous

Infrared Emitters

Most CAVEs have been using CrystalEyes infra-red emitters (or similar) [Lipton 97]. However, these emitters have both a limited range and a limited angle at which the IR signal can be received, and thus some care must be exercised during placement. Glasses that carry the synchronization signal over a wire are not usually used because all users must wear the glasses, and the extra wiring inside the CAVE would make movement more difficult.

One additional consideration is that if several people occupy the CAVE simultaneously, people may move between the IR emitters and receiver on the glasses, and the glasses will lose synchronization. This can be compensated to some extent by adding more emitters around the CAVE, so that their effective areas overlap.

Since infrared signal is just another form of light, the emitters can be placed safely behind the back-projection screens as well. However, some light (depending on the properties of the screen material) will of course be reflected, some scattered and some absorbed, so the usable range is diminished.

Most CAVE systems use between 3 and 5 emitters / wall, placed so that most of the emitters have been placed near the top of the CAVE to reduce the chance of occlusion by other users, and the rest near the screen canvas at the bottom of the screen to provide coverage when user looks downwards. The emitters are directed towards the center of the CAVE.

Tracker

In tracker placement care should be taken that the tracker itself does not intrude on the user experience. Nor should it be placed so that is it easily distracted by by the presence of magnetic materials (magnetic tracker), or the presence of other people inside the CAVE (optical tracker).

For a magnetic tracker that has a limited range, a good place to put it up in the air, or next to one of the support pylons (assuming non-distracting material), where it does not intrude on the view of the user. Sensor cabling can also be put to hang from the ceiling of the CAVE, which makes them easier to use than cabling that lies on the ground and eases the movement of the user. Figure 29, below, displays a couple of alternate tracker and wiring placements in a 4-walled CAVE, which has a front-projected floor. Placing the tracker next to the support pylon (B) is usual for a fully enclosed CAVE, whereas hanging everything from a truss (A) is common for a 4-walled CAVE.

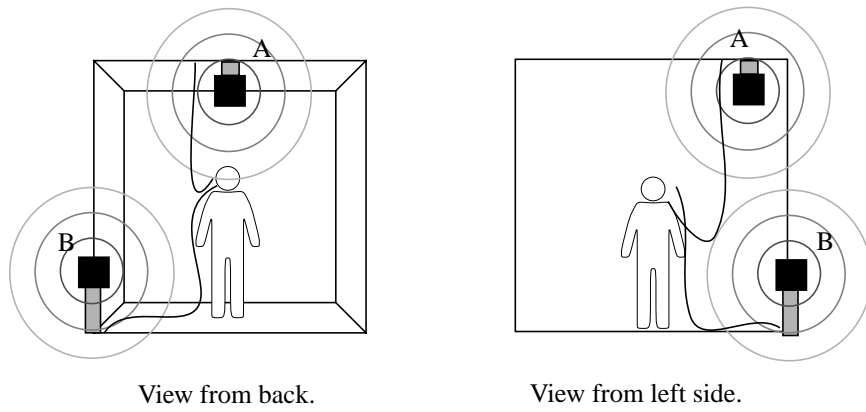


Figure 29: Alternate tracker/wiring placements.

3. HUTCave. Phase One

3.1 In the beginning...

Virtual Reality research is not a very old phenomenon in Finland [Reitmaa 95]. Unlike in rest of the Europe and US, the use of 3D graphics in CAD has not caught on until very recently, and as such, VR has not been seen as a very useful tool for industrial applications.

While there have been ideas and thoughts about teaching and researching VR within the Helsinki University of Technology [Takala 96], the situation had not been developing from its initial concept.

In May, 1997, a number of industry representatives from such companies as Konecranes, Sisu, Partek and Valmet approached the Helsinki University of Technology pronouncing their interest in VR applications and voicing their hope for seeing more VR research and education in the universities of Finland [Strandman 97]. This served as a milestone in establishing an interest and funding within the university for future research on VR and virtual prototyping.

VR research is at the writing of this paper done in at least two universities in Finland: The Helsinki University of Technology¹⁹, and the Tampere University of Technology²⁰. Applied research touching different areas of VR are known to take place in Lappeenranta University of Technology²¹, University of Jyväskylä²², University of Oulu²³, and the University of Lapland in Rovaniemi²⁴.

3.2 Project history

The idea of having a CAVE-system had long been around the HUT. With establishing of the Telecommunications and Multimedia Laboratory (TML) in 1995, the Helsinki University of Technology had finally a laboratory that could invest fully in the research of Virtual Reality.

19. <http://www.tcm.hut.fi/Tutkimus/CAVE/>

20. <http://www.dmi.tut.fi/>

21. <http://www.lut.fi/>

22. <http://www.jyu.fi/>, VR research is mostly defunct now.

23. <http://www.oulu.fi/>

24. <http://www.urova.fi/>

In July 1997 the Rector of the University finally approved a funding of approximately 1.5 Million FIM to be used in the design and construction of a CAVE unit. The initial plans were drawn quickly and negotiations to purchase the necessary equipment were started during summer, 1997. Requirements for electricity and cooling were solved temporarily by locating the Onyx2 into the Computing Centre computer room in the main building of HUT. This allowed us to utilize their cooling systems as well as the UPS power systems. The machine was installed and configured there, and later moved into the actual location of the first CAVE in the main building of HUT.

It was rather clear from the beginning that the target system was a multi-sided CAVE system, as it would allow the best immersive feeling. However, space constraints forced us to build a single-wall system at first in the basement of the main building, and later move into the new Computer Sciences building, scheduled to be finished in summer 1998. The first, single-wall CAVE is referred in this document as the “Phase I”, and the second, multi-wall CAVE “Phase II”.

Build or purchase?

One possibility that immediately occurred to us was to purchase the required CAVE system from a manufacturer (see Table 1, “Spatially Immersive Display vendors.,” on page 23). After a few offers from different vendors and discussions with international experts we decided that building a CAVE of our own would not be very difficult nor would it become too expensive. Additionally it would be a great learning experience for the laboratory.

3.3 Computing Equipment

The selection of computing equipment was done very quickly, as we already had discussed the possibilities and explored the options, which were not too many. Being the single most expensive part of the system, it could not be replaced easily, and this is why careful thought was placed on the expandability of the system.

3.3.1 Main Computer

The thought of having built a PC-based CAVE was dismissed quickly, as (in 1997) there were no systems that could’ve handled the requirements for a multiple-screen projection system (see *PC-based CAVEs* on page 57).

The logical computer choice was then the Onyx2 InfiniteReality from Silicon Graphics, Inc. The InfiniteReality2 system was not available at that time. Also, most other CAVEs have been built on the Onyx or Onyx2 platform, so compatibility with other installations would be only be limited by software.

The model chosen was the Onyx2 Rack with two R10000 CPUs, and one graphics pipeline with 2 Raster Managers with DG5-2 Display Generator, allowing the pipeline split between two graphics channels, which was an affordable solution for the single-wall HUTCAVE.

3.3.2 Auxiliary Computers

Two SGI O2 computers were purchased to function as software development platforms, and a 166MHz Pentium PC to function as the input device development platform.

The input device PC was equipped with a Voodoo 2 graphics board, Linux 2.0, and a National Instruments AT-MIO E series A/D conversion board and linked to the LAN with a 100Mb Ethernet connection. This system was then used to develop a generic input device driver system.

3.4 Software

3.4.1 Software Solutions

The commercially available CAVE systems are not necessarily shipped with any CAVE specific software, other than what normally is shipped with the graphics computer. However, most vendors have agreements with software companies and can offer deals from commercial software.

During February 1998 several of the faculty staff (Tapio Takala, Erik Bunn, Janne Jalkanen) visited some European locations to look at their software solution. During the trip we visited an EVL CAVE in Linz, Austria, and two TAN CAVEs in Fraunhofer Darmstadt and GMD/Sankt Augustin. Also, Antti Nurminen and Janne Jalkanen had previously visited an EVL CAVE in SARA Amsterdam.

Pyramid Systems CaveLib

The original Cavelib is available from Pyramid Systems Inc. freely with any purchase. Unfortunately, the cheapest item to be purchased is an ImmersaDesk, meaning that the software would simply be too expensive for an university. There is no special program for universities and the software is not available separately.

Fraunhofer's Software

Fraunhofer Research Institute IGD had developed their own tool chain for CAVE programs called Y. However, this tool chain is not based on any high-end graphics library and thus does not scale well and remained unusable for our purposes. They do, however, have a very impressive 5-sided CAVE built by TAN GmbH.

GMD AVANGO

The GMD research institute in St. Augustin, Germany, showed us their CAVE as well as the AVANGO software. We were rather impressed, and discussions were started so that the TML-laboratory could be a beta-tester for AVANGO.

Unfortunately the GMD staff was extremely unresponsive after the initial contact, and while we received the an AVANGO release later on, the software has remained unusable due to lack of documentation and communication, despite repeated attempts to get sample code or better documentation.

3.4.2 HUT Software

Some custom software was also developed at the laboratory: The LibR had been the engine behind the DIVA virtual orchestra for some time, and could be easily adapted to the single-wall CAVE display. Also a simple tool for demonstration and visualization purposes was developed locally.

LibR

LibR [Hänninen 99],[Laakso 99] is a locally developed software platform for the DIVA virtual orchestra project. It provides rudimentary support for stereo and motion tracking, but is rather limited in interactivity, and would require extensive rewriting for more complex applications other than the DIVA project itself. In addition, the rendering performance is not the best possible, since it has only a very flat scene graph structure, and is OpenGL-based. Also, it does not support any common file formats. However, it does support skeletal animation and kinematics (both forward and inverse, though the inverse IK needs to be mostly pre-calculated).

Stereofly

Since we needed a quick method of showing off the HUTCave to people, a quick hack called “stereofly” was developed [Napari 99]. It is based on IRIS Performer and is able to show stereoscopic images and provide head tracking with the Ascension MotionStar tracker. This application proved to be very versatile, though, and has served as a successful experimentation platform for a multitude of purposes. In addition, since IRIS Performer supports over 20 different file formats off-the-shelf, it has been a very useful piece of software in attempting to port 3D scenes to the HUTCave.

Unfortunately, there is no other interactivity than movement and head tracking within the Stereofly software, and thus it is not very suitable as a basis for future development. In addition, the lack of a proper wand device means that its use has been limited to “movie ride” -applications only.

3.5 Choosing the Place

Phase I was initiated in October 1997, when facilities in the basement floor of the Helsinki University of Technology main building became available, as the Laboratory of Optoelectronics moved their high-energy laser away. While the facilities were almost adequate (see Figure 30 for a picture of the facilities), it soon became apparent that fitting a 3-wall Cave into the room was impossible. Three solutions became available:

1. Remove the wall between the Cave room U027 and the neighboring room U026, which was used by the Media Brewery. This would barely allow enough room for the Cave structure.
2. Build just a single wall Cave and then get a bigger facility in the new Computer Sciences building to be completed in Summer 1998.
3. Start looking for a completely new place.

Option number three was explored briefly, but no such facilities became available in the near vicinity of the HUT. Relocation farther away was undesirable, because of the overhead involved in traveling.



Figure 30: The Media Brewery facilities where the HUTCave was located in Phase I

Removing the neighboring wall did seem to be a good idea at first, but unfortunately this would have caused some lack of space to the Media Brewery and also would have removed the operating capability of the Cave for a longer period of time. An estimated break in CAVE operations was at least three weeks, and it would have caused several problems because the wall was made out of concrete and supported a balcony, for which a new support structure should have been built.

However, we received a space in the completely new facilities would become available in the Computer Sciences building for the HUTCave, and so the research into option #2 was abandoned quickly. The subsequent reconstruction of the Cave room was then done accordingly.

3.6 Reconstruction

3.6.1 Computer Room

The original room had a full-width balcony overseeing the TV studio below, so in order to keep the noise from the computer from not disturbing the users below, a new wall was constructed. Access to the machine room was made available by adding a single door at the top of the stairs.

The double Gyproc wall was filled with rock wool, providing sufficient noise insulation. The two triple-glass windows and door were also built to keep the noise inside the machine room. We did not measure the resulting drop in sound pressure, but users did not find the remaining noise disturbing.

The floor of the balcony (as well as the operator room) was also replaced with anti-static, grounded plastic floor mat so that it would not gather static electricity which might have been harmful to the computer, and also to facilitate easier cleaning.

3.6.2 Power

The old TV studio had a lot of electricity outlets, but they were too old to meet the specifications laid out by the SGI technical staff [Zamost 98]. This posed some serious problems, since completely refitting the whole electrical system in the basement would have been too expensive.

However, the HUT computing centre was right behind the wall, and thus we were able to take the required 230V, 16A electrical outlet from there, by making a simple cabling through already existing lead-ins. While we were at it, we also laid the network cabling (Cat-5 Ethernet, 100Mb/s) through the same path.

Since the electricity from the Computing Centre was rerouted through an UPS system it had a different ground level than the electricity in the U027. This meant that monitors and other equipment had to be either insulated with opto-couplers or get their electricity from the same source.

We measured the peak current that the Onyx2, monitor, and projector took in case of a sudden power-down and the following power-up, and found that it was a safe 10A, well within the capabilities of the 16A wall socket. The audio system was de-coupled through an optical cable, so it was safe to connect to the local outlet.

3.6.3 Acoustics and lighting

The original use of the room was a TV-studio and the last users before the HUT-CAVE was the Optics laboratory, who had housed a large laser unit in the room. They had already modified the room to be completely black, and the walls were covered with acoustic material, so we could easily settle in to the new room with minimal changes to the acoustics or the lighting.

The only real problem we faced was with the control-room, which originally had a window for looking into the TV studio, but now had a solid wall. We removed a part of the wall obstructing the view and recovered the old installation, except for the glass which had been previously removed during the installation of the wall. Figure 31 shows the view from the control room to the CAVE.



Figure 31: View from the control room into the CAVE.

Unfortunately, the lights from the control room distracted the user in the CAVE, so that the control room had to be darkened during use, with only small table-top lights giving light to the operator.

3.6.4 Ventilation and cooling

The Onyx2 system dissipates a lot of heat, and removing this heat didn't at first pose a problem. The original ventilation system was designed for a TV studio and seemed quite sufficient to handle the 4 kW the Onyx2 dissipates [Zamost 98]. Since the machine room was originally an integral part of the U027 TV studio, it had no ventilation on its own and we had to redirect some of the air coming to the studio through the new wall. The outlet of air from the machine room was done through an opening above the door, which was supposed to be connected to the exit shaft.

Panic!

During winter time, no problems occurred. However, though the facility was originally designed in the 50's to have ample cooling, we found, much to our surprise, that the air conditioning equipment had actually been dismantled due to lack of use some ten years ago, and the system could only provide with air coming directly from outside, heated, unheated but not cooled. As the summer grew closer, it became apparent that the pure air ventilation was not sufficient to cool the Onyx2 system.

Heat began to be a serious problem in May with outside temperatures reaching 25° C, and we approached several specialists, and explored possibilities to see what could be done to counter the problem. Solutions were as follows:

1. **Build a completely new cooling system into the Media Brewery.** This would've benefitted the Media Brewery as well, because they had had similar problems with their video equipment during the hot summer months.
2. **Move the computer back to its original location in the Computing Centre computer room.** This was a very well cooled place, but the distance between the CAVE and the computer would have been too big for the cabling.
3. **Build a partial solution to benefit only the CAVE.** This would involve renting or buying the A/C equipment, and building a way to dispatch the extra heat.

Option 1 would have probably solved all problems in one stroke, but the University was not prepared to pay such a high price, noting that the main building would be renovated anyway in 1999, when new air conditioning would also be installed.

Going with option 2 would have meant using an optical linking system between the computer and the CAVE console and projector. This would not have been such a bad idea, and we found a suitable solution, the Video Display Extension VDE/200 unit from Lightwave communications, which transmits the monitor image, keyboard, and mouse signals through an optical cable. The maximum bandwidth for this is approximately 350 MHz to a distance of 3 km, which would have been completely sufficient for our purposes [Lightwave 98].

Options 2 and 3 were both very alluring, but when a room in the new CS building was a certainty, we no longer saw a reason in going for the expensive, but more permanent solution and leased a 4kW A/C device from a local company. It used water as a cooling agent and dumped any extra heat to a wash basin. The result was an efficient cooling system that was good enough to keep the machine room at acceptable levels at all times, even during the summer. The installation of the said device was somewhat troublesome, as the nearest water outlet was on the other side of the room, which is why temporary piping had to be installed.

Assuming that the A/C system was running at full power, the water consumption would be in the order of 0.022 l/s, or 2.6 m³/d. Using the price of water at that time, the maximum cost of cooling was 4000 FIM during the summer time (May-September), when uncooled outside air was not enough to remove the heat.

3.7 The HUTCAVE Projection System

Since it was well known that the HUTCAVE will be extended in the future to be a three-sided structure, we examined the possibility of reusing the screen and frame material. This did not prove feasible, since the framing required for a single wall is completely different from a multiple-wall CAVE, and so a simple, temporary screen was installed.

3.7.1 The Screen

The screen used was a Da-Tex Da-Lite back-projection screen, which was chosen after inquiries to several local A/V companies. The canvas was delivered with snap fasteners in leather sleeves, and the snaps had to be removed before it could be attached to the frame.

3.7.2 The Frame

The simplest and easiest construction that was free of magnetic metals that could be built to the set dimensions was wood. A local carpenter built a frame to our specifications in a few days. The screen was attached to the frame using a plastic tube inside the sleeve, which was then squeezed into an opening on the frame. See Figure 32 for details.

The frame was free-standing, and had two legs. Most of the weight was behind the screen, in order to support the frame better in case of accidental user collision with the screen or the frame. Also, we did not want the frame intrude too much into user space.

No metal was used in the construction of the frame: all joints were made with glue, which proved to be quite durable. We did not experience any problems with the frame itself, but we found that the large surface area proved to be very recep-




Figure 32: Left, attaching the screen to the frames. Right, the frame stand.

tive to air currents generated by users moving near the screen, and started to vibrate very easily. The wooden frame bent just enough to allow movements in the order of a few centimeters at the top. See Figure 33, below.



Figure 33: The frame inside the CAVE room, without projector or mirror.

3.7.3 The Projector

We evaluated and examined several possible projectors from BARCO, Electro-Home, AmPro, JVC, and Sony, but settled for the BARCO 1209s projector with P43 fast phosphor option. The BARCO projector, while the most expensive, had the best overall configurability, best resolution (2500x2000 pixels true resolution), and RGB bandwidth (120 MHz) at a decent price point. Also, it was quiet and had a decent remote controller. The BARCO projector was also a choice of other CAVE installations, and thus known to work well in such places.

The BARCO 808s projector was also a candidate, but we decided to go with the larger model because it was capable of displaying 1024x1024@114 Hz, which we wanted to do to get square pixels, even though it was slightly above the nominal bandwidth (see 2.3.3 *Bandwidth* on page 46).

3.8 Audio System

The audio system was built to be a test platform for the later multi-walled configuration. In its initial configuration, it consisted of 8 Genelec 1029A loudspeakers that are capable of producing a maximum peak acoustic output of 110 dB SPL at 1 m per pair, with a frequency response of 68 Hz - 20 KHz (3 dB). The Genelec 1029A has also been magnetically shielded, so no disturbance to or from the magnetic tracker was noticed. The built-in amplifier monitors the output levels and prevents any damage to the drivers, making the system immune to overload and spurious signals [Genelec 99].

The loudspeakers were arranged in a rough triangular mesh spanning 360 degrees horizontally, and approximately 45 degrees vertically. The upper loudspeakers were placed on 1.5 m high stands and slightly tilted towards the user, while the lower loudspeakers were placed on styrofoam bases very close to the floor. This installation was capable of providing fully spatialized in the horizontal direction sound using VBAP.

Spatialized audio was calculated on an R10K O2 workstation in the control room or the Onyx2 computer, and sent via an optical cable to a Korg 880 D/A converter, which then distributed the signals to the loudspeakers.

3.9 Installing HUTCAVE equipment

The Onyx2 graphics computer was finally installed after all of the pre-requisites were met in March 1998, the place cleaned up (dust is a killer for computers, especially dust after construction work). See Figure 34, below, for pictures on the installation phase.



Figure 34: HUTCAVE personnel installing the Onyx2 and the mirror.

Cabling was handled by making a small opening in the wall and then pulling the wires through. Especially complicated was the monitor, keyboard, and mouse cabling, since they needed to be extended to the maximum specified length (at 15 m), and there appeared some problems with “ghost” events from the mouse. However, these problems disappeared with IRIX patch #2850, so a hardware problem was not to blame. Consequently, we have been running mouse and keyboard at longer distances, with no ill effects.

Figure 35 displays the HUTCAVE environment in its Phase I operating condition.

The BARCO projector did however have some problems: initially, the BARCO factory in Belgium suffered from a fire, delaying the delivery of our projector. The local distributor loaned us a 808s -series projector, on which we could then run our installation. When the 1209s projector arrived a month late, it broke down just after a few weeks and had to be repaired.

We also found that the IRIS autoconvergence-feature on the BARCO projectors is unable to do the convergence adjustment on a back-projection screen, if the optics have been folded with a mirror. Thus convergence adjustments became a second nature, fast.



Figure 35: The HUTCAVE in Phase I configuration running an architectural application.

3.10 Miscellaneous

3.10.1 Tracking

We had acquired an Ascension MotionStar tracker previously, and this was installed into the HUTCAVE as the main tracker. We detected a huge amount of disturbance, which turned out to be one of the main power cables of the building just on the other side of the wall. Also, the floor and walls contained enough iron to make a magnetically noisy environment. Several tests were run in order to determine the least noisy sampling frequencies, and the transmitter unit was placed on a wooden pedestal at approximately 1 m from the ground, behind the user, to minimize the noise from the floor. The best sampling frequencies are available in [Napari 99].

3.10.2 Backups

Backups of the whole 31 GB of disks were done using a DLT drive. This backup device was chosen over the DAT device due to its robust construction, large capacity (40 GB compressed), and previous experiences. Incremental backups were made daily with full dumps at least every two weeks. The backup software in use is the freely available Amanda [Amanda 99].

3.10.3 Usage

Most of the usage of the Phase I HUTCAVE was testing and setting up of the system. The initial system was also used to port the DIVA virtual orchestra [DIVA 97] to an immersive environment. Also much time was used to demo the virtual wall to both students and companies. However, no real projects were started during this initial period.

Figure 36, below, shows the DIVA orchestra performing in the CAVE.



Figure 36: The DIVA Virtual Orchestra running in the Phase I HUTCAVE.

3.11 Comments on Phase I

The frame construction was rather successful for a temporary installation, but if a more permanent installation was desired, the single wall frame is not stable unless fastened from the top as well, since the large surface is very sensitive to air currents in the room, and may be distracted by the movements of users as well.

The air conditioning system worked well, and we had no further problems with overheating.

The 1209S projector worked well, but it suffered from very odd image color imbalance, and thus had to be recalibrated in BARCO factories in Belgium. During the few weeks the projector was gone, we had a 808S projector on loan from the Finnish distributor.

When the new HUTCAVE was built in the CS building, the old frame and screen were donated to the Espoo-Vantaa Institute of Technology to be used as their initial VR installation. The glued frame had to be removed in pieces, though.

For a throughout discussion of the conclusions for this phase, please see Chapter 5.: *Conclusions and Future Work* on page 102. See Figure 37, below for an image of the CAVE in operation.



Figure 37: The HUTCave phase I running.

4. HUTCAVE Phase Two

The Phase II consisted of moving the entire HUTCAVE system to the new Computer Science building. A new screen system was implemented, with three walls and a floor.

The original time frame for the Phase II was only a few months after the Phase I, but due to constant problems with the new CS building, the whole project was almost six months late, and thus we had to run with the old configuration until March 1999.

4.1 Construction work

Of course, the facilities were not immediately suitable for the HUTCAVE system. Unfortunately, we missed the deadline for construction changes by just a few weeks, and thus no major modifications could be made to the building itself. Figure 38, and Figure 39, below, display the approximate dimensions of the room. As can be seen, the ceiling is not very high, about 3.75 m, which prevents the construction of a CAVE with both the ceiling and the floor, as both of them require at least two meters free space above and below the system.

The rest of the room measures 9.51 x 10.29 m, being sufficiently big enough for a 4-sided CAVE.

4.1.1 Appearance

The original coloring of the new CAVE room was white, as it was designed to be a classroom. However, white color does not mix well with the relatively low contrast of the CRT projectors, and thus it needed to be replaced with something darker.

There are a few good choices available: your average black, dark blue, dark green, dark red, and dark brown. We looked around for examples, and settled for either dark green or dark blue, since we felt that the other colors were not very good: Black, even though it is in some ways the “optimal” color, is a very bad choice for interior design, as it is very boring, and does tend to depress people. We were seeking something calm and professional, which ruled out dark red and brown.

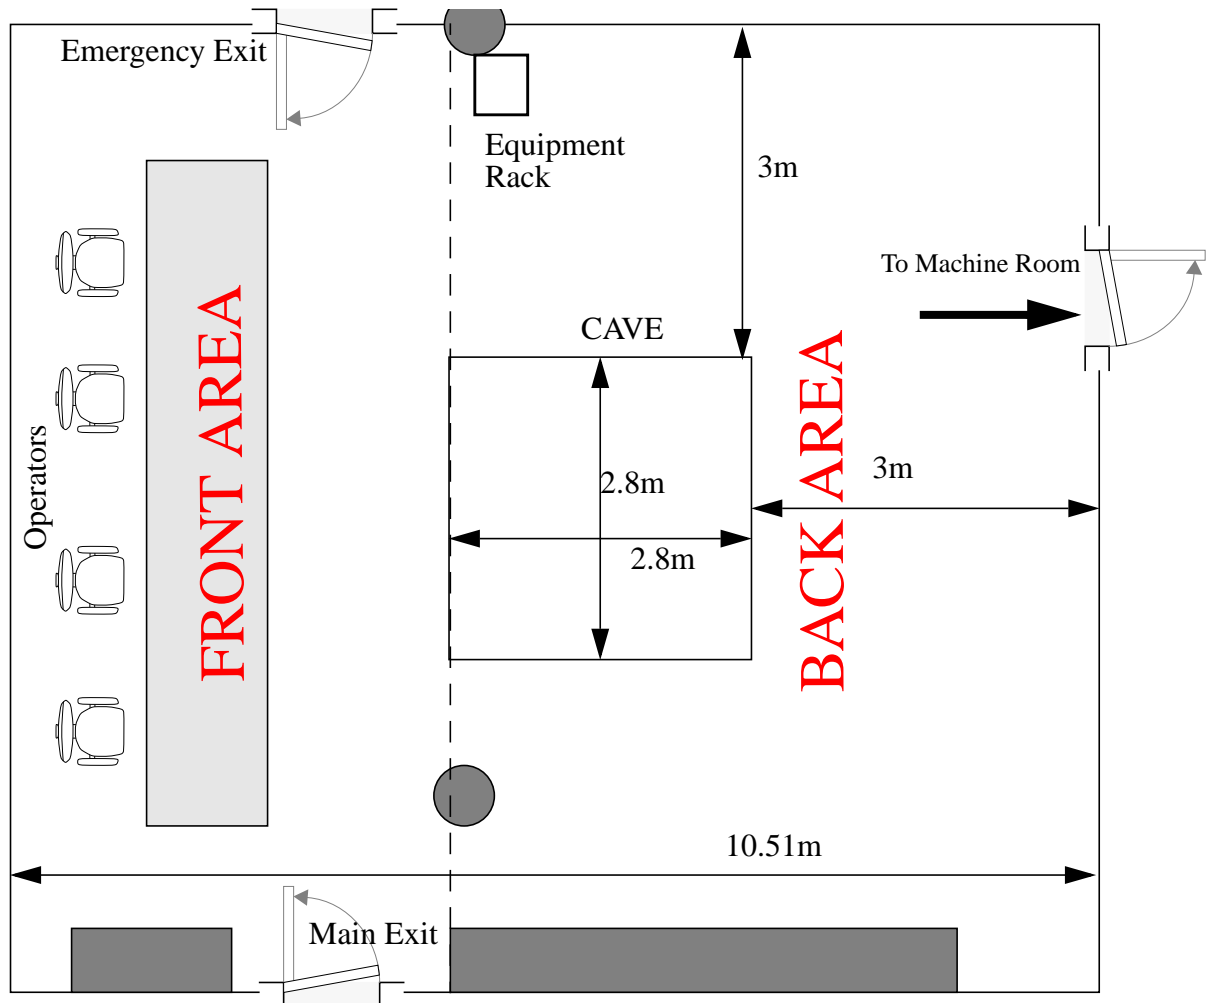


Figure 38: The CAVE room layout.

With these two options we approached the original architect of the building, who then recommended that we use dark blue with a hint of purple, so that it would go well with the overall coloring of the building. While dark green would have been more in line with the furniture available to us, we trusted the architect's words enough to order a sample piece.

The sample piece turned out to be a painted 1.5x1.5 m piece of wood. We did some visual experimentations with the sample, and were quite satisfied with the coloring, ordering the paint job.

We requested an offer from two companies to do the painting of the room, and made the choice according to those tenders. It took approximately a week to paint the room.




Figure 39: The CAVE room layout, side view.

4.1.2 Acoustics and Lighting

The walls of the room were basic computer room material, and thus very hard and echoic. Fully covering all walls with anechoic material would have been extremely expensive (there is approximately 400 m² of wall space), so only the worst flutter echoes were removed by adding absorbent material on the back wall and on some of the side walls.

The reflections between the floor and the ceiling had to be eliminated by totally covering the ceiling with acoustically absorbent material, which had the effect of lowering the ceiling by approximately 5 cm everywhere. Holes were left in all walls and the ceiling for mirrors, so that the available distance for mirrors would not be reduced. This is not acoustically a big problem, since due to the placement of the mirrors, they reflect the sound and tend to break any wave-tubes.

Figure 40 below, displays the construction of the acoustic panels inside the HUTCAVE room. The acoustical design was done by a local company specializing in the field.

The original room had very large windows, which had to be covered to make the place as dark as possible. This could have been accomplished by either covering the windows with heavy curtains or by covering them with plywood or similar. We considered curtains briefly, but we realized there would be no need for the light from the windows, and thus settled for plywood covering, since it was slightly cheaper and also possible to make 100% light-proof, unlike curtains.



Figure 40: Installing the acoustic panels inside the Cave room.

Between the windows and the wall we managed to squeeze room for “display casing”, and the 30 cm wide space was equipped with a curtain holder for future use. The doors were also covered with a soft cloth, that both killed the ambient light and provides additional acoustic dampening. See Figure 41 for an illustration.

One of our worries was the ambient noise from the nearby university cafe and storage lockers, that were embedded in the wall on the other side. However, they did not acquire that much popularity with the students, and thus noise was not a problem. Also, the acoustics of the building seemed good enough to keep noise from the cafe away. The only worry remaining was the noise caused *by* the CAVE, since the university CS library was also outside and they might be disturbed. So far, we have not received any complaints.

With ambient light gone, the lighting was produced by rows of fluorescent lamps. These lights were split into two groups: front and back. The front lamps light the operator area with the table, and the back lamps provide illumination in the “machinery” area of the room. The lamps can be turned on and off separately, and the front lamps can be dimmed down using a slider switch, which allows some light inside the CAVE (for users and operators alike) while minimizing ambient light between the projector and the screen. Real lighting for reading manuals, etc., is provided by table-top lights attached to the operator’s table.



Figure 41: Covered windows (left) and the door (right).

Two of the back area lamps had to be removed later upon the actual CAVE installation, when we found that the CAVE frame size was slightly larger than initially calculated. However, this did not significantly impact the light available.

4.1.3 Ventilation and Cooling

Since the room was originally built as a class/computer room, the ventilation was built for approximately 30 people and an equal number of computers. For the CAVE use this was more than enough, and thus there were no problems to be expected. See Table 3, below, for an approximate estimate of the required cooling power. Note that the Onyx2 was in a separate room with separate cooling system.

Table 3: Required cooling power, maximum values.

Equipment	BTU/hr.	KW
ONYX 2 (twin rack) ^a	66300	19.4
3*24" monitors	1785	0.5
4*video projector	9800	2.9
Auxiliary computers	ca.6600	ca. 2
Operators (4) ^b		0.5
Other users(approx.15) ^c		2.4

- a. The actual measured power usage of an Onyx2 single rack was approximately 7 kW. The maximum value is derived from SGI documentation.
- b. Operators, sitting down, average of 120 W.
- c. Users, moving, average of 160 W.

The worst problem was the noise from the ventilation system, which was rather disturbing. The air pumps for two other adjacent rooms reside within the CAVE room as well, and contribute to the overall noise level. The noise level was not measured, but was deemed disturbing by professionals.

Normally, the ventilation could just be turned down, but it turned out that during the first year shake-down period it must be used at full power, because of the moisture inside the building left from the construction process, so the problem stayed unsolved until the end of the year, when the ventilation became controllable again. The generic acousting process of the room (see *4.1.2 Acoustics and Lighting* on page 84) did somewhat alleviate the noise levels, though.

4.1.4 Power

Table 4 illustrates the power requirements of the different systems inside the CAVE. All values are peak powers, usually attained during startup. Maximum values have been calculated from the maximum configuration for the future (a twin rack Onyx2 with separate pipelines for each wall).

Table 4: The power requirements of the CAVE equipment.

Equipment	Real (W)	Max (W)
Onyx-2	9750	19500
Monitors	175	525
Projectors	650	1950
Auxiliary computers	3000	5000

The new building had no such problems as the old building, and we were able to draw all power from the same source, distributed by several 16A cables both the computer room and the CAVE. The power cabling was installed alongside network cabling in the ceiling, from which we used standard splitters and wire conduits to distribute power to different equipment.

A master power switch (see right) was installed for the Onyx2, as required in a fixed installation. It is also useful for emergency purposes, for example, if the air conditioning system should fail and spill its water on the floor, endangering the computer.



4.1.5 The machine room

With luck on our side, the room next to the new CAVE room was only half occupied by the Security and Administration Department, and we were able to secure approximately 10 m² of it as our computer room. The room was split in half by building a wall, and new door was constructed between the main CAVE room and the computer room to allow easy access.

During construction of the door we realized that there was a number of pipes going at approximately 20 cm from the floor on the wall. Leaving them in would have meant a very tall threshold, which in turn would've made it almost impossible to move the Onyx2 computer into the room, so they had to be removed, despite the fact that they were connected to a heating unit. Due to the way the heating unit piping was done, going around the door would have been a problem, and it was judged that with the additional heating (approximately 4x650W) provided by the projectors it is quite unlikely that the room needed any extra heating anyway. If anything, cooling would be the problem.

The room was already painted with white paint during construction work, so it was not necessary to repaint it. Lighting was installed (with two fluorescent lamps), and necessary air conditioning was also installed. As luck would have it, the room next to the computer room happened to be the main air conditioning room, so whenever more cooling was needed, it could be easily provided. Initially, only one 4kW unit was installed for cooling purposes, with room reserved for more. In order to monitor the temperature, a min-max thermometer was installed both in the machine room and the main HUTCAVE room.

At the end of June 1999, when summer was unusually hot in Finland, we experienced some problems with cooling. It turned out that the air conditioners were not installed correctly throughout the building, and we experienced a short outage in CAVE services, as condensed water started to appear on the floor of the machine room. The pump on the air conditioning machine did not work due to faulty installation, and the water flowed onto a separate spill tray reserved for these purposes. We had anticipated such a possibility, and thus had an extra backup: a secondary pump that would pump away the water from the tray (see Figure 42). Unfortunately, the tray had a hole in it, which caused some of the water to flow on the floor. Luckily, no damage was done during this brief encounter.

The analysis of the problem revealed that the spill tray had been point-welded instead of sealed, and thus the water was able to run through the seams. A temporary fix was made with silicone, which was used to seal the seams.




Figure 42: The extra pump for condensation water.

4.1.6 Other modifications

Sprinklers

Because of the placement of the CAVE unit in the middle of the room, the top mirror reflecting the image onto the floor would be placed in the almost exact center of the room. Unfortunately, there already was a sprinkler tube with a sprinkler unit right at the same exact spot. While the mirror would still have fit (after all, it is placed in an angle, leaving ample room between the back of the mirror and the ceiling), it would not have been too wise, nor legal, to leave the sprinkler beneath it.

Since the minimum legal distance between sprinkler units is four meters, it was not allowed just to remove the single sprinkler unit, but it had to be moved 1.5 meters to the side. In addition, a secondary sprinkler unit had to be installed, 1.5 meters to the other side of the original place, leaving a 3 m space where the mirror could be placed.

Because of the building contract stated that the insurance for the Computer Science building would only be valid if modifications were performed by the same operator that originally installed the sprinkler system, we could not ask for tenders. Unfortunately, this created a situation where we had to wait a considerable time until the contractor could find a suitable time slot for the re-installation, hampering the entire time schedule of the CAVE system by two weeks. We could not begin the installation of other items before the sprinkler system was in place, because of the possibility of water damage and metallic dust that is created when the tube is severed.

Finally, the modifications took approximately a day, during which half of the building was without emergency water. Of course, the fire department had to be notified, as well as the laboratory masters of the building. The original water tube was replaced with slightly larger version, because now there are four sprinklers instead of the original three.

Network

Connecting the CAVE to network did not require major modifications, since the room was already entirely equipped with CAT5 cabling, running on cable platforms hanging 50 cm from the ceiling.

The Onyx2 is connected to the local building LAN with 100 MB switched Ethernet, and the LAN is connected to the FUNET university network with a 622 MB ATM. Most of the auxiliary computers are connected similarly with a fully switched network. A 16 x 10 MB hub also exists for spurious connections, such as laptops, printers, etc.

4.2 Moving In

The original Cave was dismantled and relocated during a busy day in March 1999. The biggest problem were the two heaviest and at the same time most delicate objects: the Onyx2 computer and the Barco video projector. Also, transporting the mylar mirror was difficult, as it could be easily damaged beyond repair. Other items, mostly cabling, auxiliary computers, the tracker unit, monitors, etc. could be transported using sufficient manpower and a normal car.

The relocation was handled mostly by HUT researchers, assisted by a truck and a driver rented from a local company. The mirror was packaged inside plastic as the original packaging was thrown away due to a miscommunication just days prior to the relocation. The Onyx2 and the Barco projector were moved as-is, with minimal protection.

The relocation was quite successful, but it took weeks before the rest of the construction work was completed. Since the Onyx2 was used also as a project disk server, it was much more convenient to have in the same building as the staff, especially considering maintenance and backups, even though this meant a temporary halt in actual Cave operations. Figure 43 shows the installed Onyx2 computer inside the computer room.

During the hiatus of few weeks, most of the construction was completed, and finally in May the frame and screens were delivered.



Figure 43: The Onyx2 installed. On the right side of the Onyx, an ASCII console for controlling the computer in case of problems, and the backup devices.

4.3 HUTCAVE Projection System

4.3.1 Physical constraints

Moving from the old building to the new Computer Sciences building solved the immediate problem of tight spaces, but unfortunately the new room allocated for us was only 3.75 m high in most places.

In order to calculate the exact locations and dimensions of the mirrors and the projectors, we wrote a Matlab [Mathworks] script²⁵ which could be applied interactively against the problem. This script used ray tracing to calculate the location of the projected image through a mirror with a set of user-changeable parameters. Figure 44 below, illustrates one possible solution to placing the top mirror into the CAVE room.

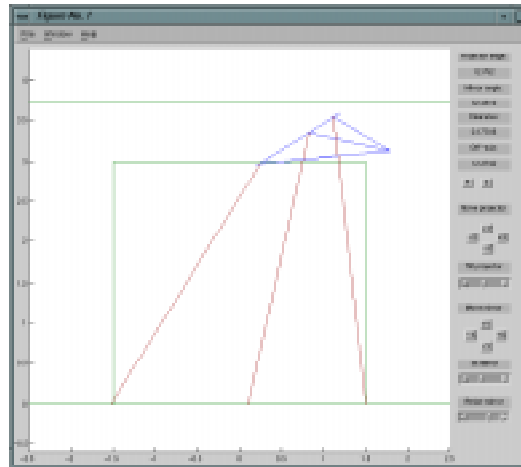


Figure 44: The Matlab projectorUI script User Interface. Image courtesy of Erik Bunn.

With the help of the aforementioned script, we were able to squeeze in the top and side mirrors within the room, but unfortunately there was very little room for error (less than 5 cm in places). To counter for this, we decided - after long debate - to make the new CAVE slightly smaller, 2800x2800x2800 mm to be exact. This gave us enough margin for errors, but did not sacrifice the immersive capabilities of the CAVE. A five-wall configuration was also considered briefly, but we decided against it due to the additional complexity, cost, and space requirements.

The original CAVE frame could in no way be installed into the new space, especially since it was glued together and could not have been removed from the premises without breaking it first. The original screen was not of very high quality either.

4.3.2 Delivery and Configuration

We queried several companies separately for the frames and the screens. Most companies were however reluctant to tackle such a complex job, and in the end only one local company was willing and able to deliver the frames and the screens. They had previous experience in building complex stage installations and back-projection systems, and had good relationships to a number of manufacturers.

25. ProjectorUI, by Erik Bunn, 1998. Script for Matlab version 5.

After the order was made, it took approximately six weeks for the equipment to arrive and to be installed, though there were serious delays in the canvas delivery.

The screen material chosen was a soft canvas, Harkness Polacoat, color grey. The frame material is unpainted aluminum. The DC magnetic tracking system already in use (Ascension MotionStar) is not affected significantly by the aluminum frame, see Chapter 2.4: *Tracking Systems* on page 50.

The CAVE edges were chosen to be welded (see *Welded Edges* on page 40) in order to minimize the visibility of the seams and edges. We also required a seamless canvas, which seriously limited the number of providers available, as very few manufacturers have the capability of producing 2.8 m wide seamless back-projection screen.

The floor was to be painted wood, and was included in the delivery.

4.3.3 Construction

The actual frame construction took only a couple of days, as the frame had been pre-manufactured and needed only assembly. See Figure 45 and Figure 46 for pictures of the different construction phases.




Figure 45: The floor and the top frame.

With the screen in place, it had to settle down a couple of days so that canvas would smooth down. After that it was re-stretched to its final form.



Figure 46: The front of the frame structure in place.

One slight problem was the future floor projector, that has to have unobstructed light path to the top mirror. The basic construction of the frame is very high (approximately 3.3 m), and there would have been no room between the ceiling and the frame for the projector. Thus the front screen had to be attached differently from the others, which were stretched with the help of a support bar and thread. The front screen is wrapped around a thick aluminum bar and attached with two-sided tape, then tightened by rotating the bar. The bar is suspended from the side frames. See Figure 47, below, for an illustration of the front screen attachment, and Figure 48 for the side screens.

Floor

The floor is made of painted wood, with a 2x2" support structure in a 40 cm grid underneath the 10 mm wood. The floor is raised from the ground by approximately 60 mm, of which 50 mm empty space is filled with rock wool to provide



Figure 47: Detail of the front corner. Notice the bar around which the front screen is wrapped.



Figure 48: The top left corner. Notice how the vertical and horizontal support bars are attached to the screen itself.

acoustic dampening. Also a small ramp has been added to the entrance of the CAVE. The floor is attached to the screens as described in 2.2.6 *Corners and Edges* on page 38. See Figure 49 on page 96, for a close-up.



Figure 49: Detail of the front left corner from the outside.

Mirrors and Projectors

The mirror and the projector could be re-used without problems from the Phase I CAVE. Unfortunately, the CAVE is slightly more elevated than the original screen, which makes the mirror stand a bit too low for our purposes, and thus the upmost few centimeters are not in use. The projector was placed on the ground and elevated slightly also. See Figure 50 on page 96.



Figure 50: The mirror from the side.

4.3.4 Problems with the projection system

Unfortunately, the single seam used by Harkness didn't prove to be very durable, and during the initial installation, the canvas was torn by its own weight. Two 10 cm rips were found later in both upper front corners. The canvas was left in place until a replacement from Harkness arrived a few weeks later, this time with double seams. The rips were deemed to be under warranty, and so no extra cost was incurred.

The frame also proved slightly problematic, as it wasn't entirely cubical in shape. This was later corrected using successive iterations and re-stretching of the canvas. The top front bar had also to be replaced to a sturdier bar, because the weight of the screen was pulling it down and the screen sagged.

The plastic clips that held the screens attached to the frame proved to break very easily, and later we had to install some metallic clips when the plastic clips began breaking down. Also, the number of plastic clips was doubled to provide less stress for each clip. Plastic clips are used to protect the screen in case an user stumbles and hits it: the clips would break down and the screen would come down in one piece.

The floor was constructed in two pieces, 1.4 m wide each, but the seam broke under the weight of a single user. This was promptly fixed by the manufacturer. The floor color has also proven to be less than optimal, since it gets stained easily. We had to establish a strict no-shoes policy for the CAVE. The paint is also of the wrong shade, and will require repainting when the floor projection is installed. It is possible that the entire floor should be covered with a front-projection screen attached with two-sided tape to the wooden surface to provide a uniform display.

The color imbalances with the projector (see 3.11 *Comments on Phase I* on page 80) returned, but not at such disturbing levels. The red tint of the image seems to occur only in certain screen modes, which stretch the resolution of the phosphor to its limits. Also, there have been some problems with the convergence of the red color, which sometimes shifts around and requires readjustment.

4.4 The audio system

The audio system is largely based on the construction on the Phase I CAVE. The same loudspeakers and other hardware are used, except that the number of loudspeakers has been increased to 16, and they have been installed in a much more spherical configuration. The loudspeakers now form a triangular mesh in three levels: one set of speakers up in the ceiling, one set at approximately head level, and one on the floor. All loudspeakers are pointed towards the epicenter of the CAVE. Figure 51 displays the different mountings of the ceiling loudspeakers, while Figure 52 shows the mountings of the floor and head-level loudspeakers.



Figure 51: Loudspeakers in the ceiling. Note the two slightly differing ways of mounting them.



Figure 52: The floor and head-level mountings of the loudspeakers.

The audio is handled via a twin-Pentium III PC running Linux 2.2, equipped with two Sonorus 16-channel sound cards. Sound is transferred via optical fibre to the Korg D/A converters in the central rack, and then to the loudspeakers using a coaxial cable. It is also possible to output the audio directly from the Onyx2 via an optical link to the loudspeakers.

The dampening properties of the screens have not been taken into account, and no measurements have so far been made. No subwoofer has been installed, and it is not known how the screens or the frame will react to high sound pressures in low frequencies.

4.5 Miscellaneous

4.5.1 Wiring

In order to keep the wiring from not disturbing maintenance too much, two cable conduits were made. The first conduit serves the wires for the upper reaches and the console (IR transmitters, upper loudspeakers, monitor, keyboard, and mouse) and the lower conduit everything else. A single route takes the lower wires next to the CAVE and then runs the rest of the wires as close to it as possible.

The upper channel uses the already existing cable platforms hanging from the ceiling and the upper trusses of the CAVE frame, which proved to be excellent wire channels. Console cabling runs from the machine room in the upper conduits, and then is taken down to the operator table via standard pylons, which also supply power. The operator table has internal cable conduits.

4.5.2 Operators

Operators are seated in front of the CAVE so that they have clear visibility into the CAVE. In total, the table has three workstations including the console and the audio/auralization PC. There is also room for two additional people next to the wall (see Figure 53, below).



Figure 53: Two operators running the CAVE tests.

4.6 Operational notes

4.6.1 Demos

Demos and demoing the HUTCAGE took a lot of time in the first few months. We found quickly that in order to minimize the impact to work, a properly working, easy to use, and well-functioning demo environment must be constructed. In addition, everybody with access to the CAVE must be trained to use it and give short presentations. Since the user interface from demo to demo tends to change, a short reference list is always useful.

For the basic demonstration architecture we used the Portalis software available from SGI. The DIVA virtual orchestra performance [DIVA 97], which had been ported over to the HUTCAGE in the early stages of development was the main showcase item, with architectural demonstrations using the Stereofly software [Napari 99] to show the immersive capabilities of the CAVE. Later, also scientific visualizations and molecular models were added.

4.6.2 Reservations

In order to keep the HUTCAGE reservations straight, we made a small Perl script that was used for making reservations over the World Wide Web. This proved to be inadequate and was replaced after a while by WebCal [WebCal 99], a Web-based calendar tool available under the GNU General Public License [FSF 91]. This tool became popular enough to be accepted by the entire laboratory staff and was rapidly adopted by other groups as well.

4.6.3 Screens and Physical Construction

While the CAVE frame construction looks very massive, it is in fact, very light - approximately 200 kg total. This means that any collisions to the support structure or the screens may cause the total misalignment of the projection system, or worse, so any movement behind the screens has been restricted.

Any ambient light from the back area is a real problem, and causes visible problems. However, if back area lighting is off, then lighting the front area is not a big problem. Of course, the installation is rather used in full darkness, but some ambient front light is acceptable, especially in demonstration use, when people need to move in and out of the CAVE constantly.

4.6.4 HUTCAGE as a working environment

The HUTCAGE proved also to be a very calm working environment for temporary employees, except for slight problems with air conditioning. While the machine room was kept at a pretty constant temperature at all times, the actual CAVE room was very cold (16-18 degrees C) during summer and very hot dur-

ing winter (in excess of 25 degrees C). The problems were probably the result of the air conditioning running at full force, since after the one-year drying period was over (*4.1.3 Ventilation and Cooling* on page 86) and the A/C could actually be turned down, the temperatures settled down. The only real complaint from the employees is that “it would be nice to see the sun every once in a while”.

5. Conclusions and Future Work

5.1 Conclusions

In this thesis, I have presented the process of building a Spatially Immersive Display, using the CAVE Automatic Virtual Environment as an example. The principles and some of the mathematics involved with the construction of an optically folded back-projection display have been explained. Screen construction and projector qualities have been discussed, as have tracking hardware and computing equipment.

Finally, two cases of CAVE construction have been presented, one single-wall display and a multi-walled display, known collectively as project HUTCAVE.

5.1.1 Quality of the work

The actual construction of the HUTCAVE is successful. The canvas is of very high quality and has almost no discernible hotspot effect. The CAVE is well within specifications, and while the edges could not be tested due to lack of projectors, the seams seem totally invisible.

The BARCO 1209S projector is very good, though rather bulky, which caused some small problems during installation and moving. It can easily display 1024x1024 pixels at 114 Hz, even though this is nominally over its bandwidth. The image quality is visually good, and scan lines are not visible in this resolution. Unfortunately the SGI monitor cannot display this resolution, and thus the work of the operator is slightly more difficult. The 808S projector which was tested prior to the 1209S could handle only 1024x768 at 120 Hz, at which resolution the scanlines were visible. The problems with the red color have been slightly annoying, but not to an unacceptable extent. The projector tends to require readjustment every few weeks, and unfortunately auto-convergence cannot be used, so this takes about an hour of maintenance monthly. This is expected to take more time, when multiple projectors are used.

The CAVE room is acoustically acceptable, though no measurements have been made. In order to make it better, the ventilation machinery would need dampening, and the acoustical properties of the mirrors and the screens would require measuring.

The final installations to the HUTCAVE occurred in September 1999. Unfortunately, not enough funding was received, and thus only a single wall is operational.

5.1.2 Organization and process

The HUTCAVE project has been mostly an exercise in process management. Some of the key lessons learned are:

- Parallelize, wherever possible. There isn't much point in waiting for one task to be finished, before starting on another - there is something you can always do.
- Plan, plan, plan. Careful planning will repay the time used ten-fold in time and cost. A bad plan is better than no plan. A good example is the paint job of the CS building CAVE: if the acoustical material had been built first, a significant amount of paint and time would have been saved, since most of the ceiling was later covered with the acoustical plates anyway.
- Never expect someone to do their job in time or correctly. Expect the unexpected, plan for the worst, and keep track of everyone's progress. Even if you become an annoyance. Know what stage your project is in and learn to express it in one sentence (because you will be asked.)
- Nothing ever happens in schedule, so plan for it. The 6-month delay in the CS building construction is a good example. We were also offered the possibility to reserve space in the new CS-II building in 1997, but as of this date (February 2000), the construction work hasn't even started...

The worst problems were people-related. Since the depression of the 90s was over, there has been lots of work to be done and too few people to do it. With HUTCAVE, the constant lack of resources slowed down the process many times, when everybody was busy with other jobs. For example, while demoing the CAVE is an important job for future projects, at one time it took simply too much time from other, equally important projects, and quickly become something that was avoided.

5.1.3 In retrospect...

Was the decision to build the HUTCAVE on our own instead of just calling any of the vendors in Table 1 on page 23 the reasonable one?

In my opinion, yes. In reality, building a CAVE of acceptable quality is rather easy - all the real issues are on the software side. Of course commercial CAVE vendors have had their act together longer and are able to provide multiple configurations at a very short notice, whereas designing and implementing everything on your own takes a lot of time. For a commercial venture, unless they are planning to build a lot of SIDs on their own, the added cost of design in man-hours makes it more sensible and cheaper to build a ready-made CAVE. For a research organization, such as HUT, where the researcher's time is cheap and not usually calculated in the costs, building a CAVE from scratch is a very viable option. During the CAVE construction a lot of knowledge about displays, projection systems, trackers, etc. is acquired, information which typically is not pub-

lished by the commercial vendors. All of this information is, however, available, but usually scattered around, and it is time-consuming to gather a complete picture about all issues in VR system design.

When the information has been gathered, though, building subsequent systems is much easier, and one can concentrate on making the quality better. Thus, if the best possible quality is required, a commercial CAVE vendor is recommended. However, if more than one CAVE is required, it is worth even for a cost-conscious business venture to think about doing the research on their own.

The biggest and most time-consuming problems encountered throughout the two HUTCAVE projects were related to the infrastructure: power management, networks, reconstruction of the rooms, colors, acoustics, and general management of things. Even if a commercial CAVE had been bought, this work would still have needed to be done. This work consisted maybe of 50% of the time.

Regardless of whether you build your own or buy a commercial one, the major problem you face is the software. As the Spatially Immersive Displays are still a relatively new phenomena, there is no real software that allows you to “just start developing”. Most of the available commercial software just supports CAVE-like displays but lack in the UI department, providing only visualization capabilities. Some recent, SID-specific software is finally coming out (such as the AVANGO [Dai 97] and VR-Juggler [Bierbaum 98]), but even they are too immature (either due to lacking hardware support or support in general) at this point to be total turn-key solutions. A research organization can afford to build their own solution, but then compatibility issues arise.

5.2 Future Work

5.2.1 Work on the HUTCAVE

Obviously, the biggest improvement in the HUTCAVE should be the acquisition of more graphics pipelines, CPUs and projectors to get a truly immersive display. Most of the work concerning the infrastructure is done - all that is required is duplication of the projection systems. If funding is acquired, then it should take approximately 10 weeks for the HUTCAVE to be fully operational.

The new projectors should be ordered, if possible, at the same time, so that the phosphors are of same age and quality. The current projector can be installed to the ceiling to provide floor projection (since it will be of slightly different color anyway due to different projection material). Also, the current installation of the projectors should be redone – they should be mounted higher on better platforms.

The mirror frames should also be replaced with adjustable frames that can reach higher.

The audio system could include more loudspeakers (24-channel audio) and at least one subwoofer for lower frequencies. The current loudspeaker array has a good frequency response between 68 Hz - 20 KHz (3 dB) [Genelec 99], so lower frequencies could use some boosting. There is also room underneath the floor reserved for additions, which might be put to use, perhaps with flat, active loudspeakers such as the EMFi [Antila 99]. Also the acoustical properties of the screens could be measured.

Work on the generic I/O PC has been progressing slowly, and it should be completed as soon as possible. The software running under Linux interfaces with a generic A/D card and provides a number of methods to handle the data from different sensors. The system has been designed to be transparent from the user's point of view and abstracts the interfaces so that any tracker can be used for any purpose.

One thing that has not yet been addressed in any great detail is the aesthetic environment of the CAVE. The user area and the machinery area should be separated with a curtain that would prevent dust, accidental light, and guests away from the delicate projection systems. It would also impact the visual appearance of the CAVE system positively. The overall user environment could also be changed by paying closer attention to decorations, etc. (One of the ideas we had consisted of placing tiny lights using optical cabling to simulate star fields all around the CAVE. The user complaint about the general darkness of the space could be alleviated somewhat.)

The long-term effects of immersion and Virtual Environments have not yet been fully investigated, though some studies suggest that immersive computer games do have effects that extend outside the game world [Mapleson 94]. It would be very interesting to conduct some studies by making people live and work inside a Virtual Environment for extended periods of time and observing them.

5.2.2 Future of Virtual Environments

A few years ago, Virtual Reality was touted as the immediate future by the media, and then dismissed when technology could not evolve fast enough to be comfortable, affordable, and useful. VR research returned to the chambers of scientists, and has been slowly leaking into the public consciousness through an alternative channel: entertainment.

The current state-of-the-art of computer gaming is heavily based on first-person, immersive experiences, with such games as DOOM, Quake, Counter-Strike, racing games, etc. Their advent boosted the 3D revolution on desktop, and very fast, affordable 3D display generators are currently available. What is high-end now, is low-end in three years.

Not only games, but also the rest of the entertainment industry with movie industry in the front lines is using the possibilities new technology is offering for immersion: New movie theatres with bigger screens and better audio are appearing, and movies are made to fully exploit these new theatres with digital surround

sound and a huge array of 3D effects that seem real, even when viewed at close-up on a 12 m screen. And people are willing to pay for all this, so the demand is there.

Virtual Environments have been accepted behind the scenes in many companies, since technologies such as Virtual Prototyping cut costs and allow faster design times, and reduce the need for costly model making. Even if actual immersive displays are not used, Computer Aided Design (CAD) tools are used to model products in 3D. Three-dimensional visualizations are also used from selling products even when they are not finished, to modeling of complex physical phenomena. The notion of teleconferencing is also very appealing to many international companies, since that would reduce the need for costly travel [Lalioti 96].

In the future, the usage of Virtual Environments is going to increase and become more and more networked. Already, the US Department of Defense is using over a billion USD by the end of 2000 to train its personnel in Distributed Interactive Simulations. The target is to be able to have 100,000 participants at the same time in the same Virtual Environment [Neyland 97].

As for individual use... SID displays are of very high quality, but they are prohibitively expensive, and require much too much room. Though, it is not impossible to imagine a TV screen the size of a living room wall, which could track the viewer with an optical system and provide a fully stereoscopic image. In fact, with current technology, it could well be built, though at a very high cost. The current trend towards home digital theatres with surround sound and large view screens seems to validate this vision of the future. With the advent of digital television and digital transmission technologies over high-bandwidth channels, this may not be such a far event.

Another thing not usually accounted for before is the emergence of mobile equipment. Current digital mobile phones offer already an almost global communication network (GSM), and new digital networks (UMTS) offer more bandwidth. Even so much, that distributing Virtual Environments to cell phones might be possible. The synthesis of mobile communication, wearable computing, and Virtual Environments is likely to change the world yet again. Maybe we will be living in a Mixed Reality in 15-20 years? (Work is already underway, see [Benford 98])

But do we want it?

Do we need it?

Do we even have a choice? Or is the all-powerful dollar (or euro) forcing these Virtual Environments down our throats?

We don't know - nobody can. However, it seems to me that if a new technology is going to be commonly adopted, it happens around 20 years after its introduction. This happened to color TV and the transistor, and is happening to digital TV as I type this. My guess would be that around 2010-2015, we will see if immersive displays and Virtual Environments are here to stay.

the focus points of the R, G, and B CRT tube, respectively. We can replace the entire 3-tube projector geometry by a single optical system that has a pseudofocus point in O, and thus calculate the AK distance with simple geometry.

The MN distance is one-half the LN distance, which is the image width on the projection screen. The angle δ and the distance FG are variables which need to be defined before calculation. The MG distance is available from the projector manufacturer for the given screen size.

The half-angle from which the image is visible from focus point O is:

$$\tan \theta = \frac{MN}{OM} \Rightarrow \theta = \text{atan} \frac{MN}{OM} \quad (\text{A-1})$$

From image geometry, we read

$$OM = MF + FG + GO \quad (\text{A-2})$$

from which it follows that:

$$FO = OM - MG + FG \quad (\text{A-3})$$

Now, we need to figure out the distance OM from the pseudofocus point O to the screen M (distances MG and FG have been determined from the projector geometry and placement). By geometry, triangles OMN and OGB are similarly shaped, and thus we get by symmetry:

$$\begin{aligned} \frac{BR}{NL} &= \frac{OG}{OM} \\ \frac{BR}{NL} &= \frac{OG}{MG + OG} \\ BR \times MG + BR \times OG &= OG \times NL \\ OG &= \frac{BR \times MG}{NL - BR} \end{aligned}$$

Similarly, the distance OM can be acquired:

$$\begin{aligned} \frac{BR}{NL} &= \frac{OM - MG}{OM} \\ OM &= \frac{MG \times NL}{NL - BR} \end{aligned}$$

By using the sine clause and noting that $\delta + \theta + \angle FKO = 180$, we get

$$\frac{\sin \angle FKO}{FO} = \frac{\sin \theta}{FK} \Rightarrow FK = \frac{FO \sin \theta}{\sin(180 - \delta - \theta)} \quad (\text{A-4})$$

and similarly, for the distance AF we get when we notice that $180 - \delta + \theta + \angle FAO = 180$:

$$\frac{\sin \angle FAO}{FO} = \frac{\sin \theta}{AF} \Rightarrow AF = \frac{FO \sin \theta}{\sin(\delta - \theta)}. \quad (\text{A-5})$$

Finally, the image width on the mirror is obviously $AK = AF + FK$.

Image Height

Let us view the situation from the side (remember, the mirror is placed upright):




Figure 55: Image geometry on the mirror, viewed from the side.

In this case, the $RV = LT$ distance measures the elevation of the projector from the ground.

The AO distance can be calculated using the sine clause from Figure 54:

$$\frac{\sin(180 - \delta)}{AO} = \frac{\sin \theta}{AF} \Leftrightarrow AO = \frac{AF \sin(180 - \delta)}{\sin \theta} \quad (\text{A-6})$$

We also get rather easily the distance $OL = LR + OR$ by

$$\frac{RG}{OR} = \sin \theta \Leftrightarrow OR = \frac{RG}{\sin \theta} \quad (\text{A-7})$$

and

$$\frac{ML}{OL} = \sin \theta \Leftrightarrow OL = \frac{ML}{\sin \theta}, \quad (\text{A-8})$$

and thus distance LR

$$LR = OL - OR = \frac{ML - RG}{\sin \theta} \quad (\text{A-9})$$

Distance AR is thus obviously

$$AR = AO - OR = AO - \frac{RG}{\sin\theta} = \frac{AF \sin(180 - \delta) - RG}{\sin\theta} \quad (\text{A-10})$$

Finally, due to geometry, we can write

$$\frac{SA}{UL} = \frac{AR}{LR} \Leftrightarrow SA = \frac{AR \times UL}{LR} \quad (\text{A-11})$$

and similarly

$$\frac{QA}{LT} = \frac{AR}{LR} \Leftrightarrow QA = \frac{AR \times LT}{LR} \quad (\text{A-12})$$

Finally, the image height is obviously

$$QS = SA + QA = \frac{AR \times UL + AR \times LT}{LR} = \frac{AR \times UT}{LR} \quad (\text{A-13})$$

Similarly the image height on the other side, from the blue CRT image can be calculated by simple substitution.

GLOSSARY

For brevity, this section only describes the concepts used often in this thesis.

Above-and-Below Stereo

A stereo format where the left-right image pairs follow each other temporally interleaved.

Autostereoscopic

A stereoscopic display which does not need any equipment from the user to see the 3D image.

Cathode Ray Tube

CRT, the standard method for image generation ever since 1950s. Used in most TV sets around the world.

Crosstalk

Phenomenon where the eye receives data that was meant for the other eye. Also known as leakage or spillage. Occurs also in audio (esp. auralization).

Field

An image frame may consist of several fields. In a progressive display one field equals one frame, and all pixels of the image are thus contained inside that field. However, in interlaced mode (such as those used in TV transmission) one frame consists of two fields, one containing the even horizontal lines, and the other containing the odd lines. These two fields are then displayed in an alternating fashion.

Field Sequential Stereo

Stereo format where you show first left eye image, then right eye image. The images are then usually then separated in the eyepiece of the user.

Frame

The basic unit of an image, consisting of one or several fields. See *Field*.

Frame Rate

1. The rate at which the display is redrawn with a new frame. If a frame consists of multiple fields, the frame rate may be lower than the actual display refresh rate. For example, the European PAL TV has a display refresh rate of 50 Hz, but an actual frame rate of 25 Hz, because one frame contains two interlaced fields.

2. The rate at which the computer can redraw the screen. For example, if a new picture is drawn every second, the frame rate is 1 fps (frames-per-second).

Head-Mounted Display

A worn display, attached to the head.

Immersion

Loss of consciousness of the outside world; feeling of being immersed into the Virtual Environment.

Interlace Stereo

A stereo format where the left-right image pair has been interleaved scan-line after scan-line.

Interlaced Display

A display mode where one frame is drawn twice, first the odd lines, then the even lines. See *Field*.

Interpupillary distance

The distance between the midpoints of the eyeballs.

Kinesthetic Feedback

Feedback relating to the use of muscles. For example weight and viscosity of a medium.

Progressive Display

A display mode where the entire image is drawn at once. See *Field*.

Scene Graph

The database representation of world data and geometry inside the computer.

Spatially Immersive Display

A display which surrounds the user, encompassing all of his visual field.

Stereopsis

The ability to interpret two separate images as the same image, with depth information.

Tactile Feedback

Stimulating the sense of touch to provide an illusion of texture, shape, temperature, etc.

Viewing Angle

The angle at which a viewer sees the target. Seeing something dead ahead means a viewing angle of zero.

Virtual Model Display

A large, three-dimensional display.

This page accidentally left blank.

REFERENCES

- [Abelson 85] Abelson H, Sussman G J, Sussman J:
Structure and Interpretation of Computer Programs
MIT Press, Cambridge, MA, USA 1985. 542 p.
- [Akeley 93] Akeley K:
Reality Engine Graphics
In Proceedings of the 20th annual conference on Computer graphics
(SIGGRAPH'96) August 2 - 6, 1993, Anaheim, CA USA.
- [Amanda 99] *The Advanced Maryland Automatic Network Disk Archiver*
University of Maryland, USA.
<http://www.amanda.org/>. Accessed on 15-Dec-1999.
- [Amiga 92] Commodore-Amiga Inc.:
Amiga ROM Kernel Reference Manual: Libraries
Addison-Wesley, USA 1992.
- [AmPro 97] AmPro Corporation:
AmPro 3600 Technical Specifications
AmPro Corporation, Melbourne, Florida, USA, 1997. Flyer.
- [Antila 99] Antila M, Nykänen H, Saarinen K:
Multichannel Electromechanical Film Panel Loudspeaker
The Proceedings of the AES 16th International Conference. Audio
Engineering Society, New York, NY 1999. pp.183-190.
- [Aoki 99] Aoki H, Schiele B, Pentland A:
Realtime Personal Positioning System for a Wearable Computers
In Digest of Papers, The Third International Symposium on Wearable
Computers (ISWC'99), October 18-19, 1999, San Fransisco, CA. IEEE
Computer Society, Los Alamitos, California, USA, 1999. pp.37-43.
- [Ascension 99] Ascension Technology Corporation
MotionStar Wireless specification
Ascension Technology Corporation, Burlington, VT, USA 1999. Available
on-line at
<http://www.ascension-tech.com/products/mswireless/mswireless.htm>
- [Auer 99] Auer T, Bratner S, Pinz A:
*The integration of optical and magnetic tracking for multi-user augmented
reality*
In Proceedings of the Eurographics Workshop in Vienna, Austria, May 31–
June 1, 1999. Springer-Verlag, Wien, Austria, 1999. pp. 43-52.
- [Azuma 98] Azuma R T, et al:
Making Augmented Reality Work Outdoors Requires Hybrid Tracking
Augmented reality: Placing artificial objects in real scenes. Proceedings of
International Workshop on Augmented Reality. AK Peters, Natick MA, USA
1999. pp. 219-224..

- [Barco 97a] BARCO Projection Systems:
BARCOGraphics 1209s Owner's Manual
BARCO nv/Projection Systems, July 1997, Belgium.
- [Barco 97b] BARCO Projection Systems:
BARCOGraphics 1209s Installation Manual
BARCO nv/Projection Systems, July 1997, Belgium.
- [Barco 98] BARCO Projection Systems:
Technical Specifications, BarcoGraphics 1209s Graphics Projector
BARCO Projection Systems, Belgium, 1998. Flyer. Available on-line at
<http://www.barco.com/projecti/pdf/bg1209s.pdf>
- [Barco 99] BARCO Projection Systems:
BARCO Elm R12 Specifications
<http://www.barco.com/projecti/products/elm/9001490s.htm>
- [Barth 00] Barth H:
Personal communication, 18-Jan-2000.
- [Benford 98] Benford S, Greenhalgh C, Reynard G, Brown C and Koleva B:
Understanding and constructing shared spaces with mixed-reality boundaries
ACM Transactions on Computer-Human Interaction. Volume 5 , Issue 3
(1998). pp. 185-223.
- [Bierbaum 98] Bierbaum A, Just C:
Software Tools for Virtual Reality Application Development
SIGGRAPH 98 Course 14 Applied Virtual Reality Notes, ACM SIGGRAPH
1998, Orlando, Florida, USA, 1998. 44 p. Also available on-line at <http://www.vrjuggler.vrac.iastate.edu/public/pubs.htm>
- [Blythe 98] Blythe D, McReynolds T, Grantham B:
Advanced Graphics Programming Techniques Using OpenGL
SIGGRAPH'98 Course. Available on-line at:
<http://www.sgi.com/software/opengl/advanced98/notes/notes.html>
- [Burdea 96] Burdea G C:
Force and Touch Feedback For Virtual Reality
John Wiley & Sons, Inc. New York, NY, USA 1996. 339p.
- [Caudell 92] Caudell T P, Mizell D W:
Augmented Reality, and Application of Heads-up Display Technology to Manual Manufacturing Processes
Proceedings of the Hawaii International Conference on System Sciences,
IEEE Press, Piscataway, NJ, USA 1992. pp.659-669.
- [Chien 99] Chien A A:
Supercomputing on Windows NT Clusters: Experience and Future Directions
In Proceedings of SC99: High Performance Networking and Computing
Conference, Portland, OR, USA, November 13 - 19, 1999. ACM/IEEE
Portland, OR, USA 1999. Available on-line at
<http://www.sc99.org/>

REFERENCES

- [Christensen 98] Christensen R R, Hollerbach J M, Xu Y, Meek S G:
Inertial force feedback for a locomotion interface.
Proc. ASME Dynamic Systems and Control Division, DSC-Vol. 64,
Anaheim, CA. Nov 15-20, 1998. pp. 119-126.
- [Clapp 87] Clapp R E:
Image Display Parameters of Dome Systems
In Proceedings of SPIE Vol. 760 Large Screen Projection Displays.
Bellingham, WA, USA 1987. pp. 22-28.
- [Cruz-Neira 95] Cruz-Neira, C:
*Virtual Reality Based on Multiple Projection Screens: The CAVE and its
Applications To Computational Science and Engineering*
Doctor's Thesis, University of Illinois, Chicago, USA 1995. 250p.
- [Dai 97] Dai P, Eckel G, Göbel M, Hasenbrink F, Lalioti V, Lechner U, Strassner J,
Tramberend H, Wesche G:
Virtual Spaces - VR Projection System Technologies and Applications.
German National Research Center for Information Technology, Institute for
Media Communication, Visualization and Media Systems Design, St August,
Germany, 1997. 74 p.
- [Darken 97] Darken R P, Cockayne W R, Carmein D:
The Omni-Directional Treadmill: A Locomotion Device for Virtual Worlds.
Proceedings of UIST'97, p.213-221. 1997.
- [Deering 92] Deering M:
High Resolution Virtual Reality
Computer Graphics, Vol 26. No. 2. July 1992. ACM SIGGRAPH, USA 1992.
pp. 195-201.
- [DIVA 97] DIVA Group: Hiipakka J, Hänninen R, Ilmonen T, Napari H, Lokki T, Savioja
L, Huopaniemi J, Karjalainen M, Tolonen T, Välimäki V, Välimäki S, Takala
T:
Virtual Orchestra Performance
In Visual Proceedings of SIGGRAPH'97, page 81, Los Angeles, CA, 1997.
ACM SIGGRAPH, USA 1997.
- [DNP] dnp denmark as:
DNP Wide Angle Projection Screens
dnp denmark as, Denmark. Available on-line at <http://www.dnp.dk/>
- [Dorfmueller 99] Dorfmueller K:
An Optical Tracking System for VR/AR Applications
In Proceedings of the Eurographics Workshop in Vienna, Austria, May 31–
June 1, 1999. Springer-Verlag, Wien, Austria, 1999. pp. 33-42.
- [Eckel 97] Eckel G:
IRIS Performer (TM) Programmer's Guide
SGI Document Number 007-1680-040. Silicon Graphics, Inc. Mountain
View, CA, USA 1997.

- [Eckel 98] Eckel G, Kempf R, Wennerberg L:
OpenGL OptimizerTM Programmer's Guide: An Open API for Large-Model Visualization
SGI Document Number 007-2852-002. Silicon Graphics, Inc, Mountain View, CA, USA 1998.
- [Eckel 99] Eckel G:
Applications of the CyberStage Spatial Sound Server
The Proceedings of the AES 16th International Conference. Audio Engineering Society, New York, NY 1999. pp.478-484.
- [Eitzmann 98] Eitzmann G:
Video Format Compiler Programmer's Guide
SGI Document Number 007-3402-003. Silicon Graphics, Inc. Mountain View, CA, USA 1998.
- [Escher 88] *The World of Escher*
Abradale Press, New York, 1988
- [Eureka 98] Eureka Magazine:
VR is having a ball
In Eureka Magazine, April 1998. Accessed on 23-Nov-1999, available on-line at:
http://www.eureka.findlay.co.uk/archive_features/Arch_Cad/vrball/vrball.htm
- [Eyre 98] Eyre J:
Fully Immersive Spherical Projection System (The CyberSphere)
On-line document, accessed on 23-Nov-1999, available at:
<http://www.ndirect.co.uk/~vr-systems/sphere1.htm>
- [Eyre 00] Eyre J:
Personal communication, 31-Jan-2000.
- [Extron] Extron Electronics:
Handbook Volume VI : Video Technologies
Extron Electronics, Anaheim, CA, USA. Available on-line at
<http://www.extron.com/>
- [Extron 99a] Extron Electronics
Handbook Volume VI: Display Technologies
Extron Electronics, Anaheim, CA, USA. Available on-line at
<http://www.extron.com/>
- [Extron 99b] Extron Electronics
Handbook Volume VI: Understanding Cables in the AV Industry
Extron Electronics, Anaheim, CA, USA. Available on-line at
<http://www.extron.com/>
- [Ferrin 91] Ferrin F J:
Survey of helmet tracking technologies
In Proceedings of SPIE Vol. 1456 Large-Screen-Projection, Avionic, and Helmet-Mounted Displays. SPIE, Bellingham, WA, USA 1991.

REFERENCES

- [Foley 87] Foley J D:
Interfaces for Advanced Computing
Scientific American, October 1987.pp. 83-90.
- [Foley 90] Foley J D, van Dam A, Feiner S, Hughes
Computer graphics : principles and practice
Addison-Wesley, Reading, MA, USA 1990. 1174 p.
- [Foxlin 98] Foxlin E, Harrington M, Pfeifer G:
ConstellationTM: A Wide-Range Wireless Motion-Tracking System for Augmented Reality and Virtual Set Applications
In Proceedings of SIGGRAPH'98, Orlando, Florida, USA, July 19-24, 1998.
Computer Graphics Proceedings, Annual Conference Series, ACM SIGGRAPH 1999.
- [FSF 91] Free Software Foundation (FSF):
GNU General Public License, version 2.
<http://www.gnu.org/copyleft/gpl.html>
Accessed on 29-Dec-1999.
- [Genelec 99] Genelec Oy:
Genelec 1029A Bi-amplified Monitoring System Data sheet
Genelec Oy, Iisalmi, Finland, 1999. Available on-line at
<http://www.genelec.com/>
Accessed on 30-Dec-1999.
- [GMO 99] Glass Mountain Optics:
Theory of Operation
<http://www.glassmountain.com/Theory.htm>. Accessed on 13-Dec-1999.
- [Golding 99] Golding A R, Lesh N:
Indoor navigation using a diverse set of cheap, wearable sensors
In Digest of Papers, The Third International Symposium on Wearable Computers (ISWC'99), October 18-19, 1999, San Fransisco, CA. IEEE Computer Society, Los Alamitos, California, USA, 1999.pp. 29-36.
- [Göbel 00] Göbel M:
GMD complements VR by olfactory display
E-mail to the info-performer@sgi.com mailing list on 23-Feb-2000.
- [Haggerty 90] Haggerty J M, Reinsch S, Bleha W P, Sterling R D:
Stereoscopic large screen displays using Liquid Crystal Light Valve projectors
In Proceedings of SPIE Vol. 1255 Large-Screen Projection Displays II. SPIE, Bellingham, WA, USA 1990.
- [Halle 97] Halle M:
Autostereoscopic displays and computer graphics
Available on-line at: <http://www.media.mit.edu/~halazar/autostereo/autostereo.html> Accessed on 11-Feb-2000.
- [Hughes 99] Hughes-JVC:
Series D300 Projector Specifications
Available on-line at <http://www.hjt.com/>

- [Huopaniemi 99] Huopaniemi J:
Virtual Acoustics and 3-D sound in multimedia signal processing
Doctor's Thesis, Helsinki University of Technology, Laboratory of Acoustics and Audio Signal Processing, Report 53. Espoo, Finland 1999. 189 p.
- [Hänninen 99] Hänninen R:
LibR - An object-oriented software library for realtime virtual reality
Licentiate Thesis. Laboratory of Telecommunications and Multimedia, Helsinki University of Technology, Espoo, Finland, 1999.
- [Ilmonen 00] Ilmonen T, Jalkanen J:
Accelerometer-Based Motion Tracking for Orchestra Conductor Following
To be published in Proceedings of 6th Eurographics Workshop on Virtual Environments, June 1 - June 2, 2000, Amsterdam, The Netherlands.
- [Inoue 90] Inoue T, Ohzu H:
Measurement of the human factors of 3-D images on a large screen.
In Proceedings of SPIE Vol. 1255 Large-Screen Projection Displays II, Bellingham, WA, USA 1990.
- [Iwata 99] Iwata H:
Feel-through: Augmented Reality with Force Feedback
In: *Mixed Reality - Merging Real and Virtual Worlds*. Ohmsa, Tokyo, Japan, 1999. 418 p.
- [Just 98] Just C, Bierbaum A, Baker A, Cruz-Neira C:
VR-Juggler: A Framework for Virtual Reality Development
In the Proceedings of 2nd International Immersive Projection Technology Workshop, Ames, Iowa, USA 1998. Available on-line at:
<http://www.vrjuggler.vrac.iastate.edu/public/pubs.htm>
- [Kalawsky 93] Kalawsky R:
The Science of Virtual Reality and Virtual Environments
Addison-Wesley, Cambridge, UK, 1993. 405 p.
- [Kato 99] Kato H, Billinghurst M:
Marker Tracking and HMD Calibration for a Video-based Augmented Reality Conferencing System
In Proceedings of 2nd IEEE and ACM International Workshop on Augmented Reality (IWAR'99), October 20-21, 1999, San Fransisco, California, USA. IEEE Computer Society, Los Alamitos, CA, USA, 1999. pp. 85-94.
- [Kaufman 74] Kaufman, L:
Sight and Mind: An introduction to visual perception
Oxford University Press, New York, USA, 1974.
- [Kawamoto 95] Kawamoto W:
3D Images That Float in Air
Byte Magazine, October 1995. Available on-line at <http://www.byte.com/>
- [Kernighan 88] Kernighan B, Richie D:
The C programming language (2nd Edition)
Prentice Hall, USA 1988. 272 p.

REFERENCES

- [Kirkpatrick 91] Kirkpatrick M D, Mihalakis G:
Projection Screens for High Definition Television
In Proceedings of SPIE Vol. 1456 Large-Screen-Projection, Avionic, and
Helmet-Mounted Displays, Bellingham, WA, USA 1991.
- [Krüger 94] Krüger W, Fröhlich B:
The Responsive Workbench
IEEE Computer Graphics and Applications, Vol 14, Issue 3, May 1994. pp.12-
15.
- [Laakso 99] Laakso M:
*More Software Tools for Virtual Reality Application Development (Maverik,
LibR, and Open Inventor)*
Internal report. Laboratory of Telecommunications and Multimedia, Helsinki
University of Technology, Espoo, Finland,1999. 22p.
- [Lalioti 96] Lalioti V, Garcia C, Frank Hasenbrink F:
Virtual meeting in cyberstage
Proceedings of the ACM Symposium on Virtual reality software and
technology 1998 November 2 - 5, 1998, Taipei Taiwan. pp. 205-212.
- [Lantz 96] Lantz E, Bryson S, Zeltzer D, Folas M T, de La Chapelle B, Bennet D:
*The Future of Virtual Reality: Head Mounted Displays Versus Spatially
Immersive Displays*
Panel, SIGGRAPH'96. Computer Graphics Proceedings, Annuals Conference
Series, New Orleans, USA 1996. pp.485-486.
- [Levin 99] Levin C:
Enter the Third Dimension
PC Magazine, December 9. 1999. Available on-line at
<http://www.zdnet.com/>. Accessed on 10-Dec-1999.
- [Lewis 71] Lewis J D, et al:
A True Three-Dimensional Display
IEEE Transactions on Electron Devices, ED-18, 9, (Sept 1971), pp. 724-732.
- [Lightwave 98] Lightwave Communications, Inc:
VDE/200 Extension Kits for Graphics Workstations
Lightwave Communications Inc, Milford, CT, USA 1998. Flyer.
- [Lipton 90] Lipton L:
Large screen electro-stereoscopic displays
In Proceedings of SPIE Vol. 1255 Large-Screen Projection Displays II,
Bellingham, WA, USA 1990.
- [Lipton 97] Lipton L:
*StereoGraphics Developers' Handbook - Background on Creating Images for
CrystalEyes and SimulEyes.*
StereoGraphics Corporation, 1997. 65 p. Available on-line at
<http://www.stereographics.com/>

- [Lipton 99] Lipton L:
StereoGraphics Stereo3D Hardware Developer's Handbook
StereoGraphics Corporation, 1999. Available on-line at
<http://www.stereographics.com/>
- [Lokki 99] Lokki T, Hiipakka J, Savioja L:
Immersive 3-D Sound Reproduction in a Virtual Room
The Proceedings of the AES 16th International Conference. Audio
Engineering Society, Inc. New York, NY, USA 1999. pp. 172-177.
- [Lucente 95] Lucente M, Galyean T A:
Rendering Interactive Holographic Images.
Proceedings of SIGGRAPH '95. In Computer Graphics Proceedings, Annual
Conference Series, 1995, ACM SIGGRAPH, New York, 1995, pp. 387-394.
- [Mann 94] Mann S:
Mediated Reality
Tech. Report TR 260, MIT Media Lab Perceptual Computing Section,
Cambridge, MA, USA, 1994.
- [Mann 98] Mann S:
*'WearCam' (The Wearable Camera): Personal Imaging Systems for long-
term use in wearable tetherless computer-mediated reality and personal
Photo/Videographic Memory Prosthesis*
Digest of Papers, Second International Symposium on Wearable Computers,
October 19-20, 1998 (ISWC'98). IEEE Computer Society, Los Alamitos, CA,
USA, 1998, pp.124-131.
- [Mapleson 94] Mapleson I C R:
Stereo 3D Vision and Realism in Computer Graphics
B.Sc. (Hons) Thesis in Computer Science, Department of Computing and
Electrical Engineering, Heriot Watt University, Edinburgh, Scotland, 1994.
Available on-line at:
<http://doomgate.gamers.org/dhs/diss/>
Accessed on 30-Dec-1999.
- [Mathworks] MathWorks, Inc:
Matlab Introduction
<http://www.mathworks.com/products/matlab/>. Accessed on 29-Dec-1999.
- [Mayer 96] Mayer T:
*Design considerations and applications for innovative display options using
projector arrays.*
Technical White Paper, presented at SPIE January 30, 1996. Available on-line
at http://www.panoramtech.com/tech_articles/Tech_Frm.htm
- [Milgram 99] Milgram P, Colquhoun H:
A Taxonomy of Real and Virtual World Display Integration
In *Mixed Reality - Merging Real and Virtual Worlds*, Ohmsa, Tokyo, Japan,
1999. 418 p.

- [Montrym 97] Montrym J S, Baum D R, Dignam D L, Migdal C J:
InfiniteReality: A Real-Time Graphics System
Proceedings of the 24th annual conference on Computer graphics & interactive techniques (SIGGRAPH'97), August 3 - 8, 1997, Los Angeles, CA USA.
- [Moore 92] Moore J R, Travis A R L, Lang S R, Castle O M:
The Implementation of a Multi-view Autostereoscopic Display
IEEE Colloquium on Stereoscopic Television, 1992. pp. 4/1-4/6.
- [Murray 97] Murray J H:
Hamlet on the Holodeck: The Future of Narrative in Cyberspace
The MIT Press, Cambridge, MA, USA, 1997. 324 p.
- [Napari 99] Napari H:
CAVE-visuaalijärjestelmä
Special Assignment in Telecommunications, Department of Electrical and Communications Engineering, Helsinki University of Technology, Espoo, Finland 1999.40 p. (in Finnish)
- [Neider 93] Neider J, Davis T, Woo M:
OpenGL programming guide: the official guide to learning OpenGL
Addison-Wesley, USA 1993. 516 p.
- [Neyland 97] Neyland D L:
Virtual Combat - A Guide to Distributed Interactive Simulation
Stackpole Books, Mechanicsburg, PA, USA, 1997. 184 p.
- [Nienstedt 93] Nienstedt W, Hänninen O, Arstila A, Björkqvist S-E:
Ihmisen fysiologia ja anatomia. (uudistettu painos.)
WSOY, Porvoo, 1993. (in Finnish)
- [Ohta 99] Ohta Y, Tamura H (editors):
Mixed Reality - Merging Real and Virtual Worlds
Ohmsa, Tokyo, Japan, 1999. 418 p.
- [Obeysekare 96] Obeysekare U, Williams C, Durbin J, Rosenblum L, Rosenberg R, Grinstein F, Ramamurthi R, Landsberg A, Sandberg W:
Virtual workbench - a non-immersive virtual environment for visualizing and interacting with 3D objects for scientific visualization
Proceedings of the conference on Visualization '96. October 28 - 29, 1996, San Francisco, CA USA. pp. 345-439.
- [Panoram] Panoram Technologies:
MegaMac 2D Visualization Computer
<http://www.panoramtech.com/MM.htm>. Accessed on 21-Dec-1999.
- [Park 99] Park J, Jiang B, Neumann U:
Vision-base Pose Computation: Robust and Accurate Augmented Reality Tracking
In Proceedings of 2nd IEEE and ACM International Workshop on Augmented Reality (IWAR'99), October 20-21, 1999, San Fransisco, California, USA. IEEE Computer Society, Los Alamitos, CA, USA, 1999. pp. 3-12.

- [Pausch 97] Pausch R, Proffitt D, Williams G:
Quantifying Immersion in Virtual Reality
Proceedings of the 24th annual conference on Computer graphics & interactive techniques. August 3 - 8, 1997, Los Angeles, CA USA.
- [Piekarski 99] Piekarski W, Gunther B, Thomas B:
Integrating Virtual and Augmented Realities in an Outdoor Application
Proceedings of the 2nd IEEE and ACM International Workshop on Augmented Reality'99. IEEE Computer Society, Los Alamitos, CA, USA 1999. pp. 45-54.
- [Pietarila 00] Pietarila J:
Image-based view registration for augmented reality
Master's Thesis. Laboratory of Telecommunications and Multimedia, Helsinki University of Technology, Espoo, Finland, 2000.
- [Pol 99] van de Pol R, Ribarsky W, Hodges L, Post F:
Interaction Techniques on the Virtual Workbench
In Proceedings of the Eurographics Workshop in Vienna, Austria, May 31–June 1, 1999. Springer-Verlag, Wien, Austria, 1999. pp. 157-168.
- [Pompei 99] Pompei F J:
The Use of Airborne Ultrasonics for Generating Audible Sound Beams
Journal of Audio Engineering Society, Vol. 47, No. 9, September 1999. pp. 726-731.
- [Pulkki 97] Pulkki V:
Virtual Sound Source Positioning Using Vector Base Amplitude Panning
Journal of Audio Engineering Society, Vol. 45, No. 6, June 1997. pp. 456-466.
- [Reitmaa 95] Reitmaa I, Vanhala J, Kauttu A, Antila M:
Virtuaaliympäristöt - kuvan sisälle vievät tekniikat
TEKES, Helsinki, Finland, 1995. 171 p. (in Finnish)
- [Rumbaugh 91] Rumbaugh J, *et al.*:
Object-Oriented Modeling and Design
Prentice-Hall, USA 1991. 500 p.
- [Sato 99] Satō K, Ōshima T, Yamamoto H, Tamura H:
Case Studies of See-Through Augmentation in Mixed Reality Project
Augmented reality: Placing artificial objects in real scenes. Proceedings of International Workshop on Augmented Reality. AK Peters, Natick MA, USA 1999. pp. 3-18.
- [Savioja 99] Savioja L:
Modeling techniques for virtual acoustics
Doctor's Thesis, Publications in Telecommunications Software and Multimedia TML-A3, Helsinki University of Technology, Espoo, Finland 1999.
- [Savioja *et al* 99] Savioja L, Huopaniemi J, Lokki T, Väänänen R:
Creating Interactive Virtual Acoustic Environments
Journal of Audio Engineering Society, Vol 47, No 9, September 1999. pp. 675-705.

REFERENCES

- [SGI 98] Silicon Graphics, Inc:
Onyx2TM Technical Report
Silicon Graphics, Inc, Mountain View, CA, USA 1998. Available on-line at <http://www.sgi.com/>
- [Sell 95] Sell M A, Sell W:
View-Master Viewers: An Illustrated History
3D Book Publications, Borger, Netherlands, 1995. 36p.
- [Sonka 93] Sonka M, Hlavac V, Boyle R
Image Processing, Analysis and Machine Vision.
Chapman Hall, Cambridge, UK, 1993.
- [Starner 98] Starner T, Schiele B, Pentland A:
Visual Contextual Awareness in Wearable Computing
Digest of Papers, Second International Symposium on Wearable Computers, October 19-20, 1998 (ISWC'98). IEEE Computer Society, Los Alamitos, CA, USA, 1998, pp.50-57.
- [State 96] State A, Hirota G, Chen D T, Garrett W F, Livingston M A:
Superior Augmented Reality Registration by Integrating Landmark Tracking and Magnetic Tracking
Proceedings of SIGGRAPH'96, New Orleans, USA 1999. pp. 429-438.
- [Strandman 97] Strandman J, Oja H, Hakkala L, Koutonen P, Ström R, Taatila S, Nurminen S:
Onko Suomi jäämässä jälkeen virtuaalisen testauksen hyödyntämisessä?
Open letter to the Rector of Helsinki University of Technology, 12.5.1997. In Finnish. 2p.
- [Stroustrup 91] Stroustrup, B:
The C++ Programming Language, 2nd Edition
Addison-Wesley, UK, 1995. 699 p.
- [Takala 96] Takala T:
Muistio: Virtuaalihuoneen rakentaminen TKK:lle
Internal Memo, 22.10.1996, Espoo, Finland. 7 p. (in Finnish).
- [Tamura 99] Tamura H, Yamamoto H, Katayama A:
Steps Toward Seamless Mixed Reality
In *Mixed Reality - Merging Real and Virtual Worlds*, Ohmsa, Tokyo, Japan, 1999. 418 p.
- [Thomas 91] Thomas M L, Reining G, Kelly G:
The Display for Advanced Research and Training: An "Inexpensive" Answer to Tactical Simulation.
In Proceedings of SPIE Vol. 1456 Large-Screen-Projection, Avionic, and Helmet-Mounted Displays. SPIE, Bellingham, WA, USA 1991.
- [Traub 67] Traub A C:
Stereoscopic Display Using Varifocal Mirror Oscillations
Applied Optics, 6,6 (June 1967), pp. 1085-1087.

- [Trayner 97] Trayner D, Orr E:
Direct View Holographic Autostereoscopic Displays
UK-VRSIG'97 Conference Proceedings. Available on-line at: <http://www.brunel.ac.uk/depts/mes/Research/Groups/vvr/vrsig97/proceed/008/hasdpape.htm>
- [Vance 87] Vance D W:
Image Distortions in Video Projection
In Proceedings of SPIE Vol. 760 Large Screen Projection Displays. SPIE, Bellingham, WA, USA 1987.
- [VRex 98] VRex, Inc:
VR-2210 High Resolution 3D Stereoscopic LCD projector
VRex, Inc, Elmsford, NY, USA 1998. Flyer, 2 p.
- [VTI Hamlin 98] VTI Hamlin Oy:
SCA600 Series Accelerometer Chip
VTI Hamlin Oy, Vantaa, Finland. 1 June 1998. Flyer, 2p.
- [Ward 92] Ward M, Azuma R, Bennett R, Gottschalk S, Fuchs H:
A Demonstrated Optical Tracked With Scalable Work Area for Head-Mounted Display Systems
In Proceedings of the 1992 symposium on Interactive 3D graphics. March 29 - April 1, 1992, Cambridge, MA USA. ACM 1992. pp 43-52.
- [Warren 98] Warren M S, Germann T C, Lomdahl P S, Beazley D M, Salmon J K:
Avalon: An Alpha/Linux cluster achieves 10 Gflops for \$150k.
In Proceedings of Supercomputing '98, Los Alamitos, 1998. IEEE Comp. Soc. 1999. Available on-line at:
<http://www.supercomp.org/sc98/>
- [WebCal 99] *WebCal - A Web Based Calendar Program*
<http://bulldog.tzo.org/webcal/webcal.html>
Accessed on 29-Dec-1999.
- [Wheatstone 38] Wheatstone C:
On some remarkable, and hitherto unobserved, phenomena of binocular vision (Part the first)
Philosophical Transactions of the Royal Society of London. Royal Society, London, UK, 1838. pp. 371-394.
- [Wu 95] Wu J-R, Ouhyoung M:
A 3D Tracking Experiment on Latency and Its Compensation Methods in Virtual Environments
In Proceedings of UIST 95, Pittsburgh, PA, USA, 1995. pp. 41-49.
- [Youngblut 96] Youngblut C, Johnston R E, Nash S H, Wienclaw R A, Will C A:
Review of Virtual Environment Interface Technology
Institute for Defense Analyses (IDA), Alexandria, VA, USA 1996. Available on-line at:
<http://www.hitl.washington.edu/scivw/scivw-ftp/publications/IDA-pdf/>

REFERENCES

[Zamost 98]

Zamost E:

Site Preparation for OriginTM Family, Onyx2TM, OCTANETM, and O2TM
Silicon Graphics Document Number 007-3452-003, Silicon Graphics Inc,
Mountain View, CA, USA 1998.

"Well, look at that. The sun's coming up."

—Sheridan in Babylon 5: "Sleeping in Light"

INDEX

A

above-and-below stereo 11
accelometer 52
accommodation 5
Alternate Realities Corporation 23
ambient light 33, 85
ambisonics 62
antialiasing 5, 56
Ascension 30, 52, 79
aspect ratio 25, 44
Augmented Reality 1
AVANGO 59
Avocado 59

B

backup 79
bandwidth 63
BARCO 24, 25, 77
BOOM 14, 51
bow 47

C

cabling 63
 power cabling 64
CAD 106
CAVE 17
 acoustics 61
 construction of 23–65
 corners 38
 dimensions 25
 edges 38
 entrance 30
 height 28
 lighting 64
 operators 64
 PC-based 57
 price of 24
 vendors 23
Cavelib 59
ccNUMA 54
CGSD 23
color matching 38
Constellation 53

contrast ratio 48
convergence 47, 48
corners 38
Critical Fusion Frequency 45
crosstalk 45
CrystalEyes 64
culling 56
CyberSphere 20

D

DataGlove 2
Da-Tex 75
demoing 100
depth
 cues 6
 perception of 6
depth cuing 7
Depth of Field 5, 12
DG5-2 55, 69
DG5-8 55
diffusion 31
Direct3D 58
disparity 7
Display Generator 55
DLT tape drive 79
Dome 19
DOOM 105
DPLEX 55
drift 52

E

eddy currents 52
edges 38
electromechanical film 62
EMFi 62
Evans&Sutherland 23
EVL 17, 37, 69
eye 5
 accommodation 5, 11

F

FakeSpace 14, 23, 24, 30, 51
fast phosphor 46

Ferry-Porter Law 45
 Field of Vision 5
 field sequential stereo 10
 Fisher-Price 8
 Flogiston 23
 floor
 construction of 37
 surface 38
 focus 56
 non-uniform 47
 foil mirrors 42
 folding optics 26
 FOV 5
 Fraunhofer Institute 70

G

G.R.A.F 23
 gain 32, 36
 Genelec 77
 geometrical distortions 47
 Geometry Engine 55
 ghosting 45
 GMD 59
 Goto Optical 23
 GPS 54
 graphics pipe 57
 gravity 52
 Grope 2

H

haptic interfaces
 DataGlove 2
 Grope 2
 Headsight 2
 Helsinki University of Technology 67
 HMD 12, 13
 Holodeck 21
 Holy Grail 4
 hot spot 34
 HP 24
 HP-UX 60
 HRTF 61
 HUT *See* Helsinki University of Technology

HUTCAVE

acoustics 84
backups 79
coloring 82
demoing 100
floor 94, 97
frame 93, 100
network 90
operators 99
Phase I 67
Phase II 82
reservations 100
screen 93
temperature 100
tracker 79

I

image separation 10, 31, 32
 IMAX 19
 ImmarsaDesk 15
 immersion 3, 105
 InfiniteReality 54
 infra-red 64
 interaxial separation *See* interpupillary distance
 InterGraph 58
 interlace stereo 10
 interpupillary distance 7
 InterSense 53
 IPD *See* interpupillary distance
 IRIS Performer 57, 58, 60
 IRIX 56, 60

J

jitter 51

K

keystone 47
 kinesthetic 4
 Kungliga Tekniska Högskolan 18, 30, 37

L

latency 51
 LEEP 2
 LibR 70

light

- ambient* 33

- spillage* 33

Lightwave Corp. 63, 74

Linux 58, 60, 69, 98

M

MAVERIK 60

MechDyne 23, 57

Media Brewery 71

Mediated Reality 1

MIPS 55

mirrors

- foil* 42

- folding optics* 26

- positioning* 26

- reflectivity* 42

- size* 28, 43

- surface glass* 42

motion parallax 7

multi-processing 57

N

network 90

non-linearity 47

O

O2 69

occlusion 6, 39

Onyx2 54, 68

- heat dissipation* 74

- picture of* 91

OpenGL 57, 58, 60

operators 64

optical tracking 53

P

Panoram 24

parallax 7

parallax displays 9

PC 57, 69, 98

Performer *See* IRIS Performer

perspective 7

phase-coherent 53

phosphor persistence 45

pincushion effect 39

pinhole camera 5, 12

Polhemus 52

power 64, 72, 87

projectors 48

- BARCO* 77

- BARCO 1209* 25, 102

- CRT* 32, 43, 46

- DLP* 49

- DMD* 32, 49

- LCD* 46, 49

- light valve* 50

- multiple* 33

Pyramid Systems 17, 23, 69

Q

Quake 105

R

Raster Manager 55, 69

real-time processes 56

refresh rate 45

registration 51

re-imaging displays 9

resolution 44

S

Scheimpflug 47

screen

- front-projection* 36, 38

- types of* 31

seagull 47

SEOS 18, 23

SGI 18, 24, 54, 58

shared memory 57

SID 12, 16, 17

side-by-side stereo 11

signal attenuation 63

signal-to-noise ratio 51

sixty-percent rule 4

skew 47

smell 4

Sogitec 24

sound 4

Spitz Inc 23

sprinklers 89

- Star Trek
The Next Generation 21
- StereoFly 70
- stereogram 8
- stereographic display 10, 31, 45
- StereoGraphics 10
- stereopsis 7, 39
- surface glass mirrors 42
- suspension of disbelief 4
- T**
- tactile feedback 4
- TAN 15, 23, 37
- taste 4
- temperature 100
- Time-of-Flight 53
- touch 4
- trackers
AC magnetic 52
acoustic 53
DC magnetic 52
mechanic 51
optical 53
placement 65
- tracking 50
- Trimension 24
- U**
- Ultimate Display 2
- update rate 51
- V**
- varifocal mirror display 9
- VBAP 62
- vibration emitters 62
- VIEW 2
- viewing angle 32
- ViewMaster 8, 9
- Virtual Environments 13, 16, 17, 51, 62, 105
future 105
moving around in 20
- Virtual Prototyping 106
- Virtual Reality
definition 1
history 2-3
in Finland 67
- VMD 12, 14
- volumetric displays 9
- VR Juggler 60
- VR Systems
classification of 12
Headsight 2
LEEP 2
Super Cockpit 3
Ultimate Display 2
VECTA 3
VIEW 2
Virtual Cockpit 3
- VRex 11, 24
- VR-Systems UK 20
- VTI Hamlin 52
- W**
- White-Line-Code 11
- Windows 57
- Windows NT 60
- X**
- X window system 54
- Y**
- Y 70
- Z**
- zero parallax setting 7