

UNCLASSIFIED

IDA PAPER P-3186

REVIEW OF VIRTUAL ENVIRONMENT
INTERFACE TECHNOLOGY

Christine Youngblut
Rob E. Johnston
Sarah H. Nash
Ruth A. Wienclaw
Craig A. Will

March 1996

Approved for public release, unlimited distribution: 16 July 1996.

19960916 139



INSTITUTE FOR DEFENSE ANALYSES
1801 N. Beauregard Street, Alexandria, Virginia 22311-1772

DTIC QUALITY INSPECTED 3

UNCLASSIFIED

IDA Log No. H 96-001239

The work was conducted under IDA's central research program. The publication of this IDA document does not indicate endorsement by the Department of Defense, nor should the contents be construed as reflecting the official position of that Agency.

© 1996 Institute for Defense Analyses, 1801 N. Beauregard Street, Alexandria, Virginia 22311-1772 • (703) 845-2000.

This material may be reproduced by or for the U.S. Government pursuant to the copyright license under the clause at DFARS 252.227-7013 (10/88).

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE March 1996		3. REPORT TYPE AND DATES COVERED Final
4. TITLE AND SUBTITLE Review of Virtual Environment Interface Technology			5. FUNDING NUMBERS IDA Central Research Program (CRP) 9001-512	
6. AUTHOR(S) Christine Youngblut, Rob E. Johnston, Sarah H. Nash, Ruth A. Wienclaw, Craig A. Will				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Institute for Defense Analyses (IDA) 1801 N. Beauregard St. Alexandria, VA 22311-1772			8. PERFORMING ORGANIZATION REPORT NUMBER IDA Paper P-3186	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Institute for Defense Analyses (IDA) 1801 N. Beauregard St. Alexandria, VA 22311-1772			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, unlimited distribution: 16 July 1996.			12b. DISTRIBUTION CODE 2A	
13. ABSTRACT (Maximum 200 words) This report reviews current interface technology for use in virtual environments. Visual, tracking, auditory, primary user input (including, for example, glove, body suit, exoskeleton, track ball, and 3-D mouse inputs), tactile, kinesthetic, full-body motion, and olfactory interface technologies are covered. In each case, the relevant human capabilities are discussed, followed by descriptions of some available commercial products and ongoing research and development efforts. This information is used as the basis for predicting how virtual environment interfaces are likely to change in the next five years.				
14. SUBJECT TERMS Virtual Environments, Virtual Reality, User Interfaces, Interface Technology.			15. NUMBER OF PAGES 282	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT SAR	

UNCLASSIFIED

IDA PAPER P-3186

**REVIEW OF VIRTUAL ENVIRONMENT
INTERFACE TECHNOLOGY**

Christine Youngblut
Robert E. Johnston
Sarah H. Nash
Ruth A. Wienclaw
Craig A. Will

March 1996

Approved for public release, unlimited distribution: 16 July 1996.



INSTITUTE FOR DEFENSE ANALYSES

IDA Central Research Project 9001-512

DTIC QUALITY INSPECTED 3

UNCLASSIFIED

PREFACE

This report presents a review of virtual environment interface technology from the perspective of the user, that is, the devices and requirements that are imposed on the user in order to interact with a virtual environment. Most of the work was performed as part of a Central Research Project, *Research and Development in Virtual Environments*, with some additional support provided by Task A-183, *Virtual Reality Assessment of the Technical Capabilities of Surgeons* and Task T-L2-1278, *Cost and Effectiveness of Multimedia Training Technologies*.

The report was reviewed by the following staff members at the Institute for Defense Analyses: Mr. Robert L. Clover, Ms. Anne A. Douville, Dr. Dexter Fletcher, Dr. Richard J. Ivanetich, Dr. Michael R. Kappel, and Ms. Julia J. Loughran.

The authors gratefully acknowledge the support provided by all those who took time from their busy schedules to discuss their products and their work.

Table of Contents

EXECUTIVE SUMMARY	ES-1
1. INTRODUCTION	1
1.1 Purpose	2
1.2 Scope	3
1.3 Limitations	3
1.3.1 Organization	4
2. VISUAL INTERFACES	5
2.1 The Human Visual System	6
2.2 Commercial Products	10
2.2.1 Datavisor Displays	12
2.2.2 FS5 Head-Mounted Display	13
2.2.3 CyberEye 100M and 100S	13
2.2.4 CyberMaxx CM1800	14
2.2.5 Dvisor Head-Mounted Display	14
2.2.6 i-glasses!	15
2.2.7 MRG Head-Mounted Displays	15
2.2.8 VIM Personal Viewer	16
2.2.9 VFX1 Head-Mounted Display System	17
2.2.10 VR4 Head-Mounted Display	18
2.2.11 VRI HMD 133	18
2.2.12 CrystalEyes Shutter Glasses	19
2.2.13 VR-1100 and VR-2000 Stereoscopic Projection Systems	20
2.2.14 BOOM 3C	21
2.2.15 Cyberface 4	22
2.2.16 Fakespace Simulation System	23
2.2.17 PUSH	23
2.2.18 Virtual Window	24
2.3 Current Research and Development	24
2.3.1 ATR Communications System Research Laboratories, Japan	25
2.3.2 British Aerospace plc, United Kingdom	27
2.3.3 BT Laboratories, United Kingdom	29

2.3.4 Canon, Inc., Japan	30
2.3.5 Dimension Technologies, Inc.	31
2.3.6 Dimensional Media Associates	33
2.3.7 IBM Thomas J. Watson Research Center and Georgia Institute of Technology	33
2.3.8 Infinity Multimedia	34
2.3.9 NASA Ames Research Center	35
2.3.10 Purdue University	36
2.3.11 Terumo Corporation, Japan	37
2.3.12 University of Illinois at Chicago	38
2.3.13 University of New Brunswick, Canada	39
2.3.14 University of Washington	40
2.3.15 Xenotech, Australia	41
2.4 Summary and Expectations	43
 3. TRACKING INTERFACES	 47
3.1 Head Tracking	48
3.1.1 Commercially Available Trackers	53
3.1.1.1 ADL-1	53
3.1.1.2 Vidtronics Wrihtrac	54
3.1.1.3 Fastrak	54
3.1.1.4 Isotrak II	54
3.1.1.5 Insidetrak	54
3.1.1.6 Ultratrak	55
3.1.1.7 Flock of Birds	55
3.1.1.8 PC/BIRD	56
3.1.1.9 SpacePad	56
3.1.1.10 CyberTrack 3.2	57
3.1.1.11 Wayfinder-VR	57
3.1.1.12 Mouse-Sense3D	57
3.1.1.13 Selcom AB, SELSPOT II	57
3.1.1.14 OPTOTRAK 3020	58
3.1.1.15 MacReflex Motion Measurement System	58
3.1.1.16 DynaSight	59
3.1.1.17 BioVision	59
3.1.1.18 Mandala Virtual Reality Systems	60
3.1.1.19 REALWare	60
3.1.1.20 RK-447 Multiple Target Tracking System	61
3.1.1.21 Head/Hand XYZ Tracker	61
3.1.1.22 GP12-3D (Freepoint 3D)	62
3.1.1.23 Logitech 3D Mouse and Head Tracker	62
3.1.1.24 MotionPak	62
3.1.1.25 GyroPoint Pro	63
3.1.2 Current R&D in Head Tracking	63

3.1.2.1	NASA Ames Research Center	63
3.1.2.2	Massachusetts Institute of Technology (MIT), Research Laboratory of Electronics	64
3.1.2.3	Computer Graphics Systems Development (CGSD) Corporation ...	64
3.1.2.4	University of North Carolina	65
3.1.2.5	Artificial Reality	66
3.1.2.6	Massachusetts Institute of Technology, Media Lab	67
3.1.2.7	Sony, Computer Science Laboratory	67
3.1.2.8	Siemens' Central Research and Development	68
3.1.2.9	Boeing Information and Support Services, CMU, Honeywell, Inc., and Virtual Vision, Inc.	68
3.1.2.10	University of Washington	70
3.2	Eye Tracking	70
3.2.1	Commercially Available Eye Trackers	70
3.2.1.1	BioMuse	70
3.2.1.2	Headhunter Head and Eye Tracking System	72
3.2.1.3	Eyegaze System	72
3.2.1.4	Dual-Purkinje-Image (DPI) Eyetracker	73
3.2.2	Current R&D in Eye Tracking	73
3.2.2.1	Hughes Training-Link Corporation	74
3.2.2.2	Interactive Systems Laboratories (INTERACT)	74
3.2.2.3	State University of New York	75
3.2.3	Summary and Expectations	76
4.	AUDITORY INTERFACES	79
4.1	The Human Auditory System	80
4.2	Commercially Available 3-D Audio Products	85
4.2.1	Acoustetron II	85
4.2.2	Protron	88
4.2.3	Q Products	88
4.2.4	RSS-10 Sound Space Processor	89
4.2.5	SDX-330 Dimensional Expander	90
4.2.6	SRV-330 Dimensional Space Reverb and SDE-330 Dimensional Space Delay	91
4.2.7	SoundStorm 3D	93
4.2.8	Virtual Audio Processing System	93
4.3	Current Research and Development	94
4.3.1	NASA Ames	94
4.3.2	Naval Postgraduate School	95
4.4	Summary and Expectations	95
5.	PRIMARY USER INPUT INTERFACES	97
5.1	Whole-Hand and Body Inputs	97

5.1.1	The Human Hand and Arm Position Sense	98
5.1.2	Commercially Available Devices	99
5.1.2.1	5th Glove	99
5.1.2.2	CyberGlove	101
5.1.2.3	Dextrous HandMaster	102
5.1.2.4	Pinch Glove	103
5.1.2.5	Position Exoskeleton ArmMaster	103
5.1.2.6	TCAS DATAWEAR	104
5.1.3	Current Research and Development	105
5.1.3.1	Armstrong Laboratory	105
5.1.3.2	Georgia Institute of Technology	106
5.2	3-D Pointing Input	107
5.2.1	Commercially Available Devices	107
5.2.1.1	CyberWand	107
5.2.1.2	Immersion PROBE-MD	108
5.2.1.3	Magellan 3D Controller and Space Controller	109
5.2.1.4	RingMouse	109
5.2.1.5	Spaceball 2003 and Space Controller	110
5.2.2	Current R&D	111
5.2.2.1	Digital Image Design Inc.	111
5.2.2.2	University of Toronto	112
5.3	Summary and Expectations	114
6.	HAPTIC INTERFACES	117
6.1	Tactile Interfaces	117
6.1.1	The Human Tactile Sense	120
6.1.2	Commercially Available Interface Devices	123
6.1.2.1	CyberTouch	125
6.1.2.2	TouchMaster	125
6.1.2.3	Tactool System	127
6.1.2.4	Displaced Temperature Sensing System	127
6.1.3	Current Research and Development	128
6.1.3.1	Armstrong Laboratory	129
6.1.3.2	Begej Corporation	130
6.1.3.3	Harvard University	130
6.1.3.4	Hokkaido University, Japan	133
6.1.3.5	Hull University, UK	134
6.1.3.6	Massachusetts Institute of Technology	134
6.1.3.7	Research Center at Karlsruhe, Germany	135
6.1.3.8	Sandia National Laboratories	136
6.1.3.9	TiNi Alloy Company	137
6.1.3.10	University of Salford, UK	138
6.2	Kinesthetic Interfaces	139
6.2.1	The Human Kinesthetic Sense	140

6.2.2	Commercially Available Devices	144
6.2.2.1	4 DOF Force Feedback Master (Surgical Simulator)	144
6.2.2.2	Force Exoskeleton ArmMaster	144
6.2.2.3	Impulse Engine Family	147
6.2.2.4	Interactor and Interactor Cushion	148
6.2.2.5	HapticMaster	148
6.2.2.6	Hand Exoskeleton Haptic Display	149
6.2.2.7	PER-Force 3DOF	150
6.2.2.8	PER-Force Handcontroller and Finger Forcer Option	151
6.2.2.9	PHANToM	153
6.2.2.10	SAFiRE	154
6.2.3	Current R&D	154
6.2.3.1	Boeing Computer Services	155
6.2.3.2	Computer Graphics Systems Development Corporation	156
6.2.3.3	Hokkaido University, Japan	157
6.2.3.4	Massachusetts Institute of Technology, Artificial Intelligence Laboratory	158
6.2.3.5	Massachusetts Institute of Technology, Department of Mechanical Engineering	159
6.2.3.6	McGill University, Canada	161
6.2.3.7	Ministry of International Trade and Industry, Agency of Industrial Science and Technology (MITI/AIST), Japan	163
6.2.3.8	Northwestern University	164
6.2.3.9	Rutgers University	166
6.2.3.10	Suzuki Motor Corporation	168
6.2.3.11	Tokyo Institute of Technology, Japan	170
6.2.3.12	University of North Carolina	172
6.2.3.13	University of Tsukuba, Japan	174
6.2.3.14	University of Washington	175
6.3	Summary and Expectations	176
7.	FULL BODY MOTION INTERFACES	181
7.1	The Human Motion Sense	182
7.2	Self-Motion Interfaces	186
7.2.1	Commercial Products	186
7.2.1.1	Aerotrim	188
7.2.1.2	CyberPak	188
7.2.1.3	CyberTron	189
7.2.1.4	DreamGlider	189
7.2.1.5	Orbotron, X-otron VR, and Supertron	190
7.2.1.6	PemRAM Motion Bases	191
7.2.1.7	SimuPod	193
7.2.1.8	SimuSled	193
7.2.2	Current Research and Development	194

7.2.2.1	Computer Graphics Systems Development Corporation	194
7.2.2.2	Cybernet Systems Corporation	195
7.2.2.3	Institute for Simulation and Training	195
7.2.2.4	Sarcos Research Corporation	196
7.2.2.5	Systran Corporation	197
7.2.2.6	University College London, UK	198
7.2.2.7	University of Tsukuba, Japan	199
7.3	Passive Motion Interfaces	201
7.3.1	Commercial Products	201
7.3.1.1	Cyber Air Base	201
7.3.1.2	Cyberchair	203
7.3.1.3	CyberMotion Interactive Motion Seat	203
7.3.1.4	IntelliSeat	204
7.3.1.5	SIM245	205
7.3.2	Current Research and Development	205
7.3.2.1	Denne Developments Limited	205
7.3.2.2	Flogiston Corporation	205
7.4	Summary and Expectations	206
8.	OLFACTORY INTERFACES	209
8.1	The Human Olfactory Sense	212
8.2	Commercial Products	213
8.2.1	BOC Group Olfactory Delivery System	213
8.2.2	Smell-Enhanced Experience System	213
8.3	Current Research and Development in Olfactory Interfaces	214
8.3.1	Artificial Reality Corporation	214
8.3.2	Marketing Aromatics, Ltd.	215
8.4	Summary and Expectations	216
9.	CONCLUSIONS	217
	REFERENCES	233
	LIST OF ACRONYMS AND ABBREVIATIONS	235
	APPENDIX A. POINTS OF CONTACT	239

List of Figures

Figure 1. LEEP Optical Viewer	10
Figure 2. Datavisor 10x/9ci.....	12
Figure 3. FS5 Head-Mounted Display	13
Figure 4. CyberEye 100M and 100S.....	13
Figure 5. CyberMaxx CM1800 Display	14
Figure 6. CyberMaxx CM1800 Tracking	14
Figure 7. Dvisor Head-Mounted Display.....	15
Figure 8. i-glasses!	16
Figure 9. MRG 2.2 Head-Mounted Display	16
Figure 10. MRG 3c Head-Mounted Display.....	17
Figure 11. MRG 4 Head-Mounted Display	17
Figure 12. VIM Personal Viewer 1000HRpv	18
Figure 13. VFX1 Head-Mounted Display System.....	18
Figure 14. VR4 and VR4000 Head-Mounted Displays	19
Figure 15. VRI HMD 133	19
Figure 16. CrystalEyes Shutter Glasses	20
Figure 17. VR-1100 Stereoscopic Projection System.....	21
Figure 18. VR-2000 Stereoscopic Projection System.....	21
Figure 19. BOOM 3C Visual Display.....	22
Figure 20. BOOM 3C Tracking	22
Figure 21. Cyberface 4.....	23
Figure 22. Fakespace Simulation System	23
Figure 23. PUSH.....	24
Figure 24. Virtual Window	25
Figure 25. HDVD Overview.....	33
Figure 26. Xenotech Autostereoscopic Display Overview.....	42

Figure 27. ADL-1.....	53
Figure 28. Wrighttrac	54
Figure 29. Fastrak	54
Figure 30. Isotrak II	54
Figure 31. Insidettrak	54
Figure 32. Ultratrak.....	55
Figure 33. Flock of Birds	55
Figure 34. PC/BIRD	56
Figure 35. SpacePad.....	56
Figure 36. CyberTrack 3.2.....	57
Figure 37. Wayfinder-VR	57
Figure 38. Mouse-Sense3D.....	57
Figure 39. SELSPOT II.....	58
Figure 40. OPTOTRAK 3020.....	58
Figure 41. MacReflex	58
Figure 42. DynaSight.....	59
Figure 43. Head/Hand XYZ Tracker	61
Figure 44. Working Volume for Freepoint Trackers	62
Figure 45. Logitech 3D Mouse	62
Figure 46. LC Technologies EyeGaze System	72
Figure 47. DPI 5.5 Eyetracker	73
Figure 48. Acoustetron II.....	85
Figure 49. RSS-10 Sound Space Processor	89
Figure 50. SDX-330 Dimensional Expander	90
Figure 51. SDE-330 Dimension Space Delay.....	92
Figure 52. 5th Glove	101
Figure 53. 22-Sensor CyberGlove	102
Figure 54. Dextrous HandMaster.....	103
Figure 55. Pinch Glove	104
Figure 56. Position Exoskeleton ArmMaster.....	104
Figure 57. Immersion PROBE-MD	109

Figure 58. Magellan 3D Controller.....	110
Figure 59. RingMouse	110
Figure 60. Spaceball 2003.....	111
Figure 61. SpaceController	111
Figure 62. Cricket	112
Figure 63. EGG.....	113
Figure 64. Human Hand Sensing Bandwidth	122
Figure 65. CyberTouch	125
Figure 66. TouchMaster.....	126
Figure 67. Tactools System.....	127
Figure 68. Displaced Temperature Sensing System	128
Figure 69. Prototype Tactile Shape Display (Harvard University).....	131
Figure 70. Temperature Display (Hokkaido University).....	134
Figure 71. Tactile Display (Sandia National Laboratories)	136
Figure 72. Programmable Tactile Array (TiNi Alloy).....	137
Figure 73. Tactile Feedback Glove (University of Salford)	139
Figure 74. Taxonomy of Manufacturing Grasps.....	143
Figure 75. 4 DOF Force Feedback Master.....	146
Figure 76. Force Exoskeleton ArmMaster.....	146
Figure 77. Impulse Engine 3000.....	147
Figure 78. Laparoscopic Impulse Engine	148
Figure 79. Interactor.....	149
Figure 80. Interactor Cushion	149
Figure 81. HapticMaster	150
Figure 82. Hand Exoskeleton Haptic Display (HEHD).....	150
Figure 83. PER-Force 3DOF	151
Figure 84. PER-Force Handcontroller	152
Figure 85. PHANToM	153
Figure 86. SAFiRE	155
Figure 87. Robotic Graphics Proof-of-Concept System Overview	156
Figure 88. Robotic Graphics Proof-of-Concept System.....	156

Figure 89. Elbow Force Feedback Display (Hokkaido University)	158
Figure 90. MSR-1 Mechanical Master/Slave.....	161
Figure 91. 7 DOF Stylus (McGill University).....	162
Figure 92. Pantograph (McGill University).....	162
Figure 93. Cartesian 6 DOF Force Feedback Manipulator (MITI/AIST)	164
Figure 94. 4 DOF Force Reflecting Manipulandum (Northwestern University).....	164
Figure 95. Second Generation Rutgers Master	166
Figure 96. Virtual Knee Palpation System	167
Figure 97. SPICE	168
Figure 98. SPIDAR.....	170
Figure 99. Molecular Docking Virtual Interface	172
Figure 100. Operation of the Atomic Force Microscope Virtual Interface	174
Figure 101. Pen-Based Force Display (University of Washington)	175
Figure 102. High Bandwidth Force Display (University of Washington).....	176
Figure 103. Aerotrim	188
Figure 104. CyberPak	188
Figure 105. CyberTron.....	189
Figure 106. DreamGlider	190
Figure 107. X-otron VR.....	190
Figure 108. Supertron	190
Figure 109. PemRAM 3 Axis Motion Base.....	192
Figure 110. PemRAM 6 Axis Motion Base.....	192
Figure 111. PemRAM Surfing Demonstration System	192
Figure 112. SimuPod	193
Figure 113. IST Treadmill Locomotion Device	195
Figure 114. UNIPORT.....	196
Figure 115. TREADPORT.....	196
Figure 116. Individual Soldier Mobility Simulator	197
Figure 117. LocoSim	198
Figure 118. LocoSim Boot Attachments	198
Figure 119. Haptic Simulator.....	200

Figure 120. Cyber Air Base	201
Figure 121. Cyberchair	203
Figure 122. CyberMotion Interactive Motion Seat.....	204
Figure 123. IntelliSeat.....	204
Figure 124. SIM245	205
Figure 125. Personal Motion Platform (PMP)	206

List of Tables

Table 1. Visual Display Device Types and Technologies	7
Table 2. Characteristics of Commercially Available Display Devices.....	11
Table 3. Tracking Technologies	49
Table 4. Characteristics of Commercially Available Tracking Devices.....	51
Table 5. Threshold Values in Free Field (MAF) and Earphone Listening (MAP).....	82
Table 6. Threshold as a Function of Subject's Age and Stimulus Frequency	83
Table 7. Characteristics of Commercially Available Auditory Products	86
Table 8. Range of Motion for Hand, Arm, and Shoulder Jointsa	98
Table 9. Characteristics of Commercial Available Glove, Body Suit, and Hand/Arm Exoskeleton Input Devices	100
Table 10. Characteristics of Commercially Available 3-D Point Input Devices	108
Table 11. Tactile Feedback Actuator Technologies	119
Table 12. Functional Features of Cutaneous Mechanoreceptorsa	120
Table 13. Stimulus Thresholds for Strain	121
Table 14. Characteristics of Commercially Available Tactile Displays.....	124
Table 15. Force Feedback Actuator Technologies	141
Table 16. Variability of Forces Exerted in Human Grasping	142
Table 17. Characteristics of Commercially Available Force Feedback Devices.....	145
Table 18. Characteristics of Commercially Available Self-Motion Products	187
Table 19. Characteristics of Commercially Available Passive-Motion Products.....	202
Table 20. Olfactory Delivery Technologies.....	210

EXECUTIVE SUMMARY

This report reviews current virtual environment (VE) interface technology, that is, the technology that allows a user to interact with a computer-generated synthetic environment. The goal for developing VE systems is to provide a user with multimodal, highly natural forms of computer interaction. Thus, the interface technology plays a critical role. By looking at over ninety available commercial products and nearly sixty ongoing R&D efforts, this report builds a picture of current interface technology capabilities and discusses how these may change in the next few years.

Why has so much recent interest focused on VE systems? Quite simply, the potential of these systems is enormous. First of all, they offer a more intuitive metaphor for human-computer interaction. The user can exploit his existing cognitive and motor skills for interacting with the world in a range of sensory modalities and, in many instances, the experience he gains in the VE is directly transferable to the real world. Also, VE technology opens up new application areas that, hitherto, have been too expensive, too dangerous, or simply impractical. While current examples of VE applications range from surgical training systems to futuristic adventure rides, the full scope of possible applications for VE systems, and their potential benefits, is still to be determined.

The technologies that are discussed are those relating to visual, auditory, tracking, primary user input (that is, glove-based, exoskeleton, joystick, trackball, 3-D mouse, and pen-based input), haptic, full-body motion, and olfactory interfaces. The role of visual interfaces is obvious and needs no discussion except to point out that humans are strongly oriented to their visual sense, even to the extent of giving precedence to the visual system if there are conflicting inputs from different sensory modalities. While tracking is a type of interface that is largely transparent to the user, it is critical in keeping the VE system informed about user movements so that sensory inputs can be correlated to the user's position. Auditory interfaces can play a key role in providing informational inputs to the user, increasing the realism of a simulated environment, and promoting a user's sense of presence in a VE. In addition, they are used in sensory substitution where, for example, a tone is sounded to indicate when a user comes "in contact" with a virtual object and so substitute for the sense of touch. The term primary user input interfaces is used here to refer to those means whereby the user provides direct input into the VE system, for example, commands that control the operation of the system. Haptic interfaces provide the tactile and kinesthetic feedback arising from user contact with objects in the environment. Full-body motion interfaces fall into two categories. Active self-motion interfaces allow a user to move freely

through an environment, for example, walking over various types of surfaces or climbing stairs as necessary. Passive motion interfaces reflect the use of some type of vehicle to move a user through the environment. The final interface technology to be discussed is that of olfaction, where odors are used to provide the user with additional sensory cues about his environment.

At the present time, visual, tracking, and primary user input interfaces are the ones best suited for practical VE applications. In each of these cases, there is a solid basis of commercial products for potential users to choose from. Auditory interface technology is on the verge of becoming ready for use in practical applications. Indeed, increased use of auditory interfaces is the major change anticipated in VE interfaces in the next couple of years. Haptic interface technology still is largely in the research domain. Although various haptic feedback devices have been developed and a few have been used in prototype applications, the only practical use of haptic interfaces that is expected to occur within the next two to three years is with devices that are purpose-built for highly specialized applications. Widescale use of this technology is unlikely within the next five years. With respect to full-body motion interfaces, there are several entertainment systems that support limited types of highly specialized movement. Support for more general types of active user movement is exclusively a research topic with a variety of different approaches being investigated. The next few years likely will see continuing work of this type, perhaps with some prototype applications being developed. Active motion interfaces are not expected to become suitable for general use within the next five to seven years. Current work on interfaces for passive motion is focusing on a new breed of motion chairs, which will probably become widely used by the entertainment market in the near future. Olfactory interface technology is the least mature of all the technologies discussed here and unlikely to see practical usage within the three to five year timeframe.

All current VE interface technologies suffer from some limitations, even the more mature visual, tracking, and primary user input technologies. In no instance does the interface technology match human capabilities for the relevant sensory modality.

In the case of visual interfaces, head-mounted displays (HMDs) are the primary means of achieving an encompassing visual volume. HMDs suffer from several problems, with the most serious limitations being:

- Inadequate display update rates when responding to user head movements.
- Inability to provide both high resolution and a broad field of view.
- Weight that imposes an inertial burden and low levels of comfort that prevent prolonged use.

All these problems are well recognized and the first two are likely to be substantially reduced in the next few years with advances in liquid crystal diode (LCD) technologies. While smaller, lighter weight displays will help to reduce overall HMD weight, the necessity for bulky optics means that weight will continue to be a problem. A former problem,

the expense of commercial HMDs, is becoming less serious as more low cost devices are becoming available, although these require the user to make some compromises in resolution and/or field of view.

So far passive glasses have not been widely used in VE applications, although new microelectronic fabrication techniques for creating polarizing filters at the pixel level may change this trend. Shutter glasses are quite widely used, usually with cathode ray tube (CRT) or projection displays. Here again, advances in LCD technology are likely to see an impact as LCD displays with faster switching time will help in reducing crosstalk problems. There is much research and development in the area of autostereoscopic displays and a small number of products is likely to come to market in the next two to three years. Retinal displays are a new topic of research and development. While they have the potential for providing a fully encompassing visual display without the weight and limited resolution and field of view of current HMDs, it will likely be some years before these become available for practical use.

Systems for tracking head, hand, and body movements are available and many have seen widespread use. Even so, low latency, high accuracy systems for tracking in noisy, unprepared environments do not exist. The most serious shortcoming of current technology is the following:

- Inherent limitations in some combination of accuracy, intrinsic latencies, working volume, susceptibility to interference of obscuration, and cost.

Again, these are well-recognized problems that are expected to be the focus of near-term research and progress, especially for magnetic trackers, is expected. The most significant improvements in tracking performance, however, are expected to come from the use of hybrid trackers where many of the limitations inherent in a particular technology can be overcome. Only limited research is being performed on wide-area trackers and this type of tracking interface is not expected to see widespread use any time soon.

Eye tracking also is a less mature type of tracking technology. The major problems appear to relate to accuracy and intolerance to user head movements. The increased use of multimodal interfaces (in both VE and non-VE applications) that can benefit from the ability to monitor the direction of the user's gaze, however, is opening up new potential markets that should encourage further development of this type of interface technology.

A number of 3-D sound processors that can be used in VEs are commercially available. These range in capability from systems for use with PCs, to high-end professional audio systems. However, a number of questions need to be answered and further research done before virtual audio can become a practical tool. Serious limitations are the following:

- Inability to represent sounds as being located in front of the user and to adjust sound spatialization to head movements.
- Inadequacies in acoustic signal generation.

Near term work is expected to focus on these areas, continuing to improve the realism and full-surround capabilities of the technology. Crucial support for this work will come from the development of improved algorithms, based on a more thorough understanding of how humans perceive sounds. As digital signal processing becomes less expensive, virtual audio is likely to become more widespread; it is expected to become a common component of VE systems within the next five years.

The development of glove-based devices for user input is an area of current growth. The set of available products do allow the use of natural hand gestures for certain, limited interactions with a VE but the primarily shortcoming remains:

- Limited joint resolution and poor discrimination between gestures.

While improvements in sensor technology might help reduce this problem, it is likely that advances in software-based gesture recognition will play a more important role. Gloves already are a fairly common VE input device but their use is expected to become more widespread as gesture recognition capabilities improve. There seems to be little ongoing research looking at the use of exoskeleton-based devices and these are not expected to be widely used, but limited to highly specialized applications.

A fairly diverse range of 3-D mouse-based, joystick, trackball, and pen-based input devices is available. These products represent mature technology and, while new products may appear over time, no major changes in this area are expected.

Tactile and force feedback interfaces for VEs have been able to exploit previous work in the areas of, respectively, sensor substitution devices for the disabled and teleoperation. Both represent active areas of research and development. In the case of tactile interfaces, researchers are investigating how to provide contact force, slip, texture, vibration, and thermal sensations. Products intended to simulate contact forces that occur when a user touches a virtual object and products that provide temperature feedback are already commercially available. The ability to support other types of tactile sensation is more problematic. In addition to shortcomings in tactile interface hardware, much work is still needed in developing the software models needed to drive the generation of tactile signals. The major limitations in the area of tactile feedback can be summarized as follows:

- Limitations in the ability to represent surface characteristics such as texture, local shape, and slip.
- Inability of devices to present a range of tactile sensations.
- Limitation of tactile feedback to small areas.
- Lack of models and algorithms for efficient generation of tactile signals.

As stated, this is an active area of research and much progress is expected over the next few years. Nevertheless, although several prototype applications are expected, tactile interfaces are unlikely to see common use within the next two to three years. Practical applications should start appearing shortly thereafter.

The majority of current force feedback devices can be distinguished as exoskeleton devices that deliver forces to the shoulder, arm, or hand; tool-based devices that deliver forces to the hand via a knob, joystick, or pen-like object held by the user; thimble-based devices that deliver forces to the user's fingertips; or robotic graphics systems that move real objects into place to provide natural forces to the user. Each type of device is limited in the type of interactions it can support. Consequently, although several devices are on the market, each provides very different capabilities and is suitable for different types of application. The serious limitations of force feedback interfaces are, in many respects, similar to those given for tactile interfaces:

- Inability to provide force feedback for a variety of different VE interactions.
- Limitation of force feedback to a restricted number of joints.
- Intrusive nature of force feedback devices and their constraints of user movement.
- Lack of common models and algorithms for efficient generation of kinesthetic signals.

This too is an active area of research where technology advances can be expected to occur in the next five years. It is likely, however, that initial advances will be application-specific, largely in the area of medical applications where there is much interest in supporting the simulation of surgical procedures. Only a couple of the current devices have seen any practical use and more widespread use is not foreseen in the next few years.

A number of approaches and devices have been developed to facilitate a user "moving" through a VE. The simplest, and most common of these, is for the user to point in the desired direction and for the visual scenes to be adjusted accordingly. A number of entertainment systems provide highly specialized interface devices allowing, for example, the user to simulate hang gliding or sledding. Unfortunately, there has been little progress in providing more general interfaces that allow a user to simply walk or run through a VE. Technology that can support a user moving through a large area or across a surface with varying characteristics has only recently begun to be investigated. A number of diverse designs for interface systems have been proposed and a few prototypes built, using both mechanical and non-mechanical approaches. While such systems may see use as advanced prototypes, none are expected to come into common practical use within the next three to five years. The potentially large entertainment market also has fostered the development of passive motion interfaces. In the last year, several motion chairs have been developed that employ techniques ranging from inflatable chair cushions to motion bases in order to provide the user with a sense of motion. These devices may become widely used for a diverse range of low-cost simulators.

Attention only recently has turned to providing olfactory cues for VEs. There are a few commercial systems available, but none of these is capable of full control of the user's breathing space. Some prototype systems that do provide such control, at least one of which

is intended to be portable, are being developed. Nevertheless, the demand for olfactory interfaces is relatively small and this technology is expected to mature slowly and not become practically available in the near future.

In addition to further research and development on actual interface hardware and software, all the areas of interface technology discussed in this report will benefit from a better understanding of the role of sensory cues and human perceptual issues. This improved understanding not only is required to know how sensory cues can be delivered or simulated, but when and how they should be used. This is not to say that full fidelity of sensory cues is the ultimate goal. Even if achievable, high levels of fidelity would be expensive and not always desirable. What is needed is to determine the fidelity required for specific applications and how best to satisfy those requirements.

1. INTRODUCTION

This report reviews virtual environment (VE) interface technology from the perspective of the user, that is, the devices and requirements that are imposed on the user in order to interact with a VE. In this report the term VE is used synonymously with *virtual reality* and *synthetic environment*. There is no widely accepted definition of the term, and the approach chosen here is to describe a VE system as a computer-generated world with which the user can interact with the purpose of altering the state of the user or of the computer (Durlach and Mavor, 1995). The intent is to provide the user with a meaningful environment with which he can interact in a natural, multi-modal manner. For example, in a medical training application, a surgeon can practice particular surgical procedures on a virtual patient. In addition to visual images, the surgeon's major form of interaction with the system is by means of specially-modified versions of his customary instruments that provide realistic haptic feedback sensations as the surgeon manipulates virtual body tissues (Hunter et al, 1993). A virtual prototyping application might surround a designer with the visual representation of a new Space Station design which he could then move through to determine the ease of access to critical maintenance hatches. In this case, the major form of interaction would arise through the user's body movements, not only in walking to different parts of the space craft, but in seeing whether he could reach a given bolt with enough maneuvering space to exert the necessary torque to release it (Tanner, 1993).

Both of the above examples are representative of immersive VE systems, where the user is essentially surrounded by the virtual world to the exclusion of the real world. VE systems may also be non-immersive. In this case, the user views the virtual world indirectly through a computer monitor or some other display and, typically, interacts with the VE using more traditional keyboard, mouse, and trackball interfaces. A third alternative is augmented reality systems where the virtual world is superimposed over the real world. Here the intent is to supplement the real world with useful information, for example, guidance in performing a real world task. This report focuses on interface technology for immersive VE systems, although some of the material also is applicable to non-immersive and augmented reality systems.

Why has so much recent interest focused on VE systems? Quite simply, the potential of these systems is enormous. First of all, they offer a more intuitive metaphor for human-computer interaction. The user can exploit his existing cognitive and motor skills for interacting with the world in a range of sensory modalities and, in many instances, the experience he gains in the VE is directly transferable to the real world. Also, VE technology

opens up new application areas that, hitherto, have been too expensive, too dangerous, or simply impractical. The examples already given illustrate cases where previously unavailable training and practice opportunities can be provided without risk to actual patients, and how critical design decisions can be checked early in the design process without the construction of expensive physical mock-ups. A VE system can also be used to simulate a world not based on reality, or a world distorted in some meaningful way. For example, the ability to manipulate the laws of gravity while observing the effect on objects offers a valuable tool for high school physics education (Dede, Loftin, and Regian, 1994). Research chemists can benefit from a VE system that allows them to directly manipulate representations of binding forces between molecules (Brooks et al, 1990). The full scope of possible applications for VE systems, and their potential benefits, is still to be determined.

In some respects, VE systems are not new. Aircraft simulators have been in use by the Department of Defense (DoD) and the airline industry for many years, and are an obvious example of what are now called VE systems. In general terms, any computer simulation is itself a VE, although the user interaction with such simulations historically has been very restrictive. The primary innovation in today's VE systems lies in the user interface, that is, the ability to support multi-modal interaction with a simulation.

Before continuing, it is important to note that VE technology is still in its infancy. Instances of all the mentioned examples already exist or are under development and there have been VE systems that have demonstrated practical effectiveness (see, for example, (Hancock, 1993), (Magee, 1995), and (Finch et al, 1995)). Nonetheless, current systems are quite primitive, particularly with respect to their user interfaces. Not only are advances in interface hardware and software required, but a better understanding of many user issues is needed. Without question, VE technology is promising, but one that has yet to fully mature.

1.1 Purpose

One of the major recent publications in this field is the National Research Council's review of VE scientific and technological challenges (Durlach and Mavor, 1995). Prepared at the request of a consortium of federal government agencies, this review provides a overview of the current state of research and technology, a summary of major applications areas, and recommendations intended to guide a "rational and systematic development" of the field. Recent books, most notably those by Burdea and Coiffet (1994) and Barfield and Furness (1995), also provide excellent descriptions of the current state of VE technology and how these technologies work. The purpose of the current work is not to repeat any of these previous efforts, but to provide a supplement to that work. Focusing exclusively on VE interface technology, it describes some currently available commercial products and some current research and development efforts. This information provides a baseline against which the current state of art can be extrapolated to predict how VE interface technology might evolve over the next few years. In addition, the descriptions of available products are expected to provide a useful resource to potential consumers, while the descriptions of on-

going research and development efforts might serve to help researchers keep abreast of the overall directions of current work.

1.2 Scope

The interface technologies that are discussed are visual, auditory, tracking, primary user input (that is, glove-based, exoskeleton, joystick, trackball, 3-D mouse, and similar device-based input), haptic, full-body motion, and olfactory interfaces. The role of visual interfaces is obvious and needs no discussion except to point out that humans are strongly oriented to their visual sense, even to the extent of giving precedence to the visual system if there are conflicting inputs from different sensory modalities. While tracking is a type of interface that is largely transparent to the user, it is critical in keeping the VE system informed about user movements so that sensory inputs can be correlated to the user's position. Auditory interfaces can play a key role in providing informational inputs to the user, increasing the realism of a simulated environment and promoting a user's sense of presence in a VE. In addition, they are used in sensory substitution where, for example, a tone is sounded to indicate when a user comes "in contact" with a virtual object and so substitute for the sense of touch. The term primary user input interfaces is used here to refer to those means whereby the user provides direct input into the VE system, for example, commands that control the operation of the system. Haptic interfaces provide the tactile and kinesthetic feedback arising from user contact with objects in the environment. Full-body motion interfaces fall into two categories. Active self-motion interfaces allow a user to move freely through an environment, for example, walking over various types of surfaces or climbing stairs as necessary. Passive motion interfaces reflect the use of some type of vehicle to move a user through the environment. The final interface technology to be discussed is that of olfaction, where odors are used to provide the user with additional sensory cues about his environment.

The scope of the work reported here was limited by available resources. Invariably, the choice was made to focus on technologies that are specific to VEs at the expense of those that are well-defined areas in their own right. Accordingly, speech recognition and generation, natural language processing, gesture recognition, computer image generation, and cabin simulator technologies are not covered here. Similarly, application-specific interface devices, such as special-purpose knobs and switches, or steering wheels, are excluded.

With respect to research and development efforts, the focus is on *ongoing* work. No attempt is made to provide complete references to earlier efforts, although some mention of previous work is made where this directly impacts the current work discussed.

1.3 Limitations

The commercial products and research efforts discussed in this report were identified from a number of sources. In the case of products, the primary sources were published lists of vendors and advertisements found in several of the trade magazines. Research efforts were primarily identified from researchers already known to be active in the VE field

and the technical literature. The resulting information should not be regarded as comprehensive, but rather as providing a representative sampling of the available products and ongoing research. In particular, work that is regarded as proprietary, or unpublished for any reason, was unlikely to be identified.

The VE field is an active and fast-moving one and, therefore, the information contained in this report potentially has a short half-life. In each case, the product and research descriptions have been reviewed by the applicable vendor or researchers¹. Accordingly, the details reported were accurate prior to the release of this report but subsequently may have changed. Ideally, the same details would be provided for each commercial product of the same type. While attempts were made to preserve consistency wherever possible, in some case the desired information was unavailable.

1.4 Organization

The following sections of this report discuss each of identified types of interface technology in turn. Where applicable, these discussions start with an overview of capabilities of the relevant human sensory systems. This material not only builds a picture of the psychophysical interactions that take place, but indicates some requirements for particular interface devices. Short descriptions of commercially available products are followed by descriptions of on-going research in the area. The discussion of each technology area is concluded with a summary discussion and some statements of expectations for technology advances in the next few years. The report closes with a concluding section that provides an overall picture of the major limitations in current VE interface technology and gives some projections on how this technology is expected to advance in the near future.

¹ All specification data was supplied by the vendors and researchers concerned, and was not subject to independent analysis.

2. VISUAL INTERFACES

A VE imposes a number of requirements for visual displays. The most significant of these are stereoscopic vision and the ability to track head movements and continually update the visual display to reflect the user's movement through the environment. In addition, the user should be surrounded by visual stimuli of adequate resolution, in full color, with adequate brightness, and high-quality motion representations. These requirements are extremely demanding, given the capabilities of current displays and computing platforms, although progress is rapidly being made.

Another major challenge is the provision of display hardware that is not only capable of providing the necessary quality at an acceptable cost, but that minimizes the impact on the user. Currently available displays for immersive VEs typically require the user to wear a head-mounted display (HMD) or some form of special glasses. These introduce a range of new issues, such as ergonomics and health concerns, which are particularly critical in the case of HMDs.

There are a very large number of different techniques for providing stereoscopic vision. These fall into the general categories of HMDs, active (shutter) glasses, passive glasses, and autostereoscopic displays. Present-day HMDs use a technique in which each eye is provided with a separate display, together with optics that magnify the image and allow the user to focus at some depth other than the surface of the display screens. The displays and associated optics are mounted in a helmet type device, often with a position tracker and headphones attached. A new version of an HMD, now in the research stages, is based on the retinal display, in which images are not displayed on a screen, but are created by directing a beam of light (such as from a laser) to the retina of the eye.

A similar, but less encumbering alternative is the use of special glasses. In active glasses, electronic shutters are mounted in the place of the lenses of eyeglasses and, hence, these devices are often called shutter glasses. The shutters are monochrome LCDs that are used to display an opaque image to one eye and a transparent image to the other, continually switching between eyes. The user looks at a cathode ray tube (CRT) monitor or projection screen that shows left and right images as sequential fields, and that also generates a synchronization signal (such as from an infrared emitter) that controls the timing of the shutters. Passive glasses, on the other hand, use an approach in which perspective views for each eye are encoded in the form of either color (for example, red for one eye and green for the other) or polarization of the light, with the "lens" for each eye containing a filter that passes only the appropriate image intended for each eye.

Autostereoscopic displays do not require the user to wear any form of display or special glasses, although a head tracker may be needed. A variety of techniques are used. In some systems, lenses behind or in front of a display screen focus the image so that each eye necessarily sees a different image. In other systems, barriers such as vertical bar in front of the display prevent both eyes from seeing the same image. Another approach uses beams of light to scan a 3-D volume that serves as a projection screen, with the beams reflected to display a pixel at a given coordinate.

The relative strengths and weaknesses of all the different types of devices, and the accompanying technologies, are summarized in Table 1.

Relatively little is known about the conditions that provide a sense of immersion in a virtual environment. While stereoscopic vision is generally considered to be necessary for true 3-D vision and a sense of immersion, the capability of changing the visual image in response to head movement (as occurs in a real environment) may be more critical for 3-D vision than stereopsis in some conditions. For example, in one experiment (Ware, Arthur, and Booth, 1993), subjects made more errors in a task requiring 3-D vision when a stereoscopic display was used without head-coupling than when a head-coupled monoscopic display was used. (The fewest errors of all were made with a head-coupled stereoscopic display.)

This section continues with an introduction to the human visual system that presents those aspects of the visual sense that have an impact on display requirements. Commercially available display devices, intended for use with VEs, are then described, followed by descriptions of on-going R&D in this field. The section finishes with a summary of the current status of technology in VE visual displays and gives some projections for expected advances in the next few years.

2.1 The Human Visual System

The human visual system is very complex and only partially understood. It is clearly powerful with a very high bandwidth and remarkable ability to resolve detail, color, texture, and depth. The visual system also involves substantial capabilities for processing information and complex networks of neurons in both the eye and brain are devoted to visual processing (Hubel, 1963). Vision is generally considered the most dominant sense, and there is evidence that human cognition is oriented around vision, with people often using visual imagery as mediating representations for thought (Kosslyn, 1980; 1994). Thus, it is natural for high-quality visual representations to be considered critical for VEs.

The visual system consists of the eyes, certain pathways and intermediate processing centers that carry visual information from the eyes to the brain, and the visual cortex of the brain. Light enters the eye through the cornea, a transparent bulge, and some proportion of the incoming light passes through the pupil, a circular opening that is similar in form and

Table 1. Visual Display Device Types and Technologies

Technology	Description	Strengths	Weakness
HMDs	Use dual monitors (CRT or LCD) and special optics to present a different image to each eye. The monitors and lenses are mounted in a helmet-type device, usually with a position tracker and ear phones attached.	- Encompassing visual volume - Relative freedom of movement - (CRT-based): Small, high-resolution, high luminance, monochrome displays - (LCD-based): Color displays with low voltages near user's head	- Weight and inertial burden - Single viewer - Bulky optics that introduce distortions - (CRT-based): High voltages near user's head, mechanical or electrical color filtering techniques needed for color, sources of distortion - (LCD-based): Low resolution, slow switching time
Active Glasses	Special glasses with electronic shutters that display images to each eye alternately. Require a special monitor or projection screen that presents left and right eye images as sequential fields.	- Low weight - Multiple viewers possible (static stereo scenes for all but head-tracked user)	- Interocular crosstalk - Not suited for encompassing visual volume - Effectively halves frame rate
Passive Glasses	Special glasses with filters that pass only the image intended for each eye. Require a special monitor or screen in which perspective views for each eye are encoded in form of color or light polarization.	- Low weight - Multiple viewers possible (static stereo scenes for all but head-tracked user) - Inexpensive glasses (though projection display may be expensive)	- Not suited for encompassing visual volume - Poor contrast - Color coding results in eye fatigue
BOOM	Use dual monitors (CRT or LCD) and special optics to present a different image to each eye. The monitors and lenses are mounted in device suspended from a boom in front of the user.	- Weight counter-balanced - Low latency - Ease of switch to keyboard operation	- User movement restricted by mechanical linkages
Autostereoscopic Displays	Common approaches use either lenses positioned behind or before a display screen, or physical barriers in front of the display to cause each eye to see a different image.	- Unencumbering - Multiple viewers possible - Ease of switch to keyboard operation	- Low resolution - Slow switching time - User movement restricted to limited viewing area - Not suited for encompassing visual volume

function to the aperture of a camera. Muscles in the middle of the iris (the colored part of the eye) contract to increase or decrease the size of the pupil. Light that passes through the pupil enters the crystalline lens, a transparent structure that has muscles surrounding it that can rapidly alter its shape, allowing the eye to focus on particular objects, a process known as accommodation. Images are refracted by the lens and projected onto the retina, a thin layer of neural tissue that makes up most of the eye's interior (Davson, 1989).

The retina is often thought of as analogous to the sensor in a television camera that converts light to electricity and does perform this function, but it is also a very complex visual information processing system whose function goes far beyond creating electrical impulses. The structure of the retina is complex, and consists of several layers, with one layer devoted to photoreceptors, and others to concentrating and processing the output of the photoreceptors. The photoreceptor layer itself is relatively complex: there are two main types of photoreceptors, rods, of which there are approximately 120 million, and cones, of which there are approximately 8 million. Rods are used primarily for night vision, have poor sensitivity to detail, and are not sensitive to color, though they are extremely sensitive to low levels of light. Cones have good resolution for detail and are sensitive to color. In fact, there are three different kinds of cones, known as blue cones (with a peak sensitivity at 435 nm), red cones (peak at 565 nm), and green cones (peak at 535 nm). The rods and cones are not at all equally distributed in the retina: the fovea, an area of the retina upon which the central image is projected, has a very heavy concentration of cones and very few rods, while the periphery, or remainder of the field, has a heavy concentration of rods but few cones, with the density of cones decreasing with the distance away from the center of the fovea. As will be discussed later, it is possible to make use of this in the design of visual displays that economically display information in color with high resolution in the fovea and in black and white at low resolution in the periphery.

Complex circuitry in the retina, the lateral geniculate nucleus (a structure between the eye and the brain that does preprocessing), and the visual cortex of the brain perform a variety of processing. Some of this is concerned with color, while other circuitry is concerned with shape. In particular, there are layers of neural tissue that process information so as to identify increasingly abstract information. Thus, lower-level layers detect edges (with some neurons sensitive to horizontal edges, for example, and others vertical), and higher level layers detect more abstract shapes, such as curves that make up objects. Human vision also is highly sensitive to both depth and motion perception. The visual system uses a complex variety of information to determine the depth of an object, such as binocular disparity and linear perspective cues.

The field of view is the angle that an eye, or pair of eyes, can see in either the horizontal or vertical dimension. The total horizontal field of vision of both human eyes is about 180° without eye movement or, allowing for eye movements to the left or right, the total field of vision possible without moving the head is 270°. The vertical field of vision is typically over 120°. While the total field is not necessary for a user to feel immersed in a visual

environment, there is a belief among some in the community that at least 90°, and perhaps 110°, is necessary for the horizontal field of vision.

Visual acuity is the ability of the eye to resolve two stimuli separated in space. This measure is significant in that it has implications for image resolution: it is desirable for resolution to be sufficiently high that the ability of the eye to resolve stimuli, rather than the resolution of an image being displayed, is the limiting factor. Visual acuity depends significantly on both luminance levels and whether the stimuli is presented in the fovea or the periphery, with a difference of more than 20:1 between the high acuity seen with bright light in the fovea and the poor acuity resulting from dimly lit stimuli presented in the periphery (Mandelbaum and Sloan, 1947). In general, this reflects the much greater visual acuity for cone cells as opposed to rods. Assessments of visual acuity vary substantially depending upon whether they are determined by calculating the size of the retinal receptors or experimentally by psychophysical measurements, with results ranging from 0.5 to 20 seconds of arc (Davson, 1989). It is common, however, in discussions of visual acuity and its implications for display resolution, to use the more conservative figure of 30 seconds of arc for the smallest resolution visible (McKenna and Zeltzer, 1992).

Visual simulations that work by rapid successive presentations of images to the eye—as in the case of motion pictures, television, or computer-controlled displays—should preferably have successive frames presented at or above a certain rate. This rate is the critical fusion frequency, the point at which stimuli are perceived as a continuous stimulation (as fused) rather than distinct successive images (Davson, 1989). In general, the greater the luminous intensity of a stimuli, the higher the frequency at which successive images must be presented to avoid flicker. In the fovea, the critical fusion frequency is generally proportional to the logarithm of the luminance of the stimuli over a wide range (0.5 to 10,000 trolands). At high luminances, the critical fusion frequency is about 50-60 Hz, while at very low luminances it may be as low as 5 Hz. In addition, the critical fusion frequency is proportional to the size of the area of the retina in which the image falls, as well as other factors (Landis, 1954). While flicker is undesirable—it is annoying, makes perception more difficult and presumably disturbs the sensation of immersion—there is typically a trade-off needed between image complexity and susceptibility to flicker in systems with fixed computational power, and in some applications it may be preferable to tolerate flicker at least some of the time to gain increased scene complexity. Under most conditions, a 60 Hz refresh frequency (used for television in the United States), will result in an absence of flicker. According to McKenna and Zeltzer (1992), a rule of thumb in the computer graphics industry suggests that below about 10-15 Hz, objects will not appear to be in continuous motion, resulting in distraction.

The human eye is sensitive to an extremely wide range of light levels, about 12 logarithmic units. About 6 of these levels are under rod vision, while the other 6 are under cone vision. However, the eye cannot operate at any given time across this entire range: instead, the eye adapts to a given level of light, largely by mechanisms involving the light-sensitive chemicals in the receptor neurons in the retina. Such adaptation is very rapid when light lev-

els increase, but take on the order of minutes or tens of minutes when light levels decrease. For a certain state of adaptation, the eye is sensitive to about two orders of magnitude of brightness.

2.2 Commercial Products

The vast majority of current commercial products are HMDs. Such displays range from high-end, expensive products such as the Cyberface 4, Datavisor, Ericsson HMD, Fakespace Simulation System, and Stereoviewer 1, which have prices around \$40,000 to \$55,000, to medium-price systems such as the CyberEye, MRG Head-Mounted Displays, and Private Eye, which have prices above \$1,000 but less than \$10,000, to quite inexpensive consumer product displays costing less than \$1,000, including the i-glasses!, and VFX1. Other products described include CrystalEyes shutter glasses, the Virtual Window autostereoscopic display, and the VR -1100 and VR-2000 passive glasses projection systems. The major characteristics of many of these products are summarized in Table 2.

Before describing the commercially available products, it is useful to briefly discuss the special optical system used in most HMDs. Marketed by Leep Systems, Inc., the LEEP (Large Expanse Extra Perspective) Optical Viewer (also known as product ARV-1) is a unique set of lenses that has become a de facto standard for stereoscopic vision optics in the VE industry. It is a lens system that magnifies the images of the LCD or CRT displays of a stereoscopic viewer, so as to increase the field of view of the display. A photograph of the LEEP Optical Viewer is given in Figure 1.

The product is based on a design originally intended for use in stereoscopic photography (Howlett, 1983). At that time, stereoscopic color photography was practical only with very narrow (20° to 50°) fields of view and the LEEP invention made possible wide angle view capture for slides that must be an eye-spacing apart. The strong positive distortion needed to record a wide angle for each eye introduces an additional lateral

chromatism, in which optics refract a beam of light at a different angle if it is of a different wavelength. The result of this, in the case of optics such as the LEEP, is differential magnification depending upon wavelength, which is observed as red and blue fringes at the edges of a field. The design of the LEEP lenses deliberately ignored the problem of lateral chromatism. The camera and stereoscopic viewer both used the same complementing chromatism distortion lenses for original photography and for viewing. The chromatic aberration and distortion of the image stored on the film was corrected by an inverse aberration distortion effect when the stereograph was viewed. The LEEP optics result in a wide angle "fish-eye" lens system that, as a whole, does not distort the position of objects in the image.




Photo courtesy of Leep Systems, Inc.

Figure 1. LEEP Optical Viewer

Table 2. Characteristics of Commercially Available Display Devices

Technology	Product	Vendor	Resolution (in pixels)	Field of View	Weight	Price
HMD	CRT-based					
	Datavisor 10x/9ci	n-Vision, Inc.	1280 x 1024	80° horiz	4.2 lb	Contact vendor
LCD-based	Ericsson HMD	Ericsson Microwave Systems	1280 x 960	55° horiz, 41° vert	2.3 kg	Approx. \$55,000
	FS5	Virtual Research Systems, Inc.	Not applicable	55° - 76° diag	38 oz	\$19,900
	CyberEye 100M and 100S	General Reality Company	420 x 230	22.5° horiz, 16.8° vert	14 oz	\$1,995, \$2,495
	CyberMaxx CM1800	VictorMaxx Technologies	320 x 400	58° horiz	14 oz	\$889
	Dvisor HMD	Division, Inc.	345 x 259	105° horiz, 41° vert	8 lb	~\$7,000
	i-glasses	Virtual I/O	640 x 480	30° diag	8.5 oz	\$599-\$799
	MKG 2.2 HMD	Liquid Image Corp.	240 x 720	84° horiz, 65° vert	4 lb	\$3,495
	MKG 3c HMD	Liquid Image Corp.	768 x 556	84° horiz, 65° vert	4 lb	\$5,500
	MKG 2.2 HMD	Liquid Image Corp.	480 x 234	61° horiz, 46° vert	2.5 lb	\$2,195
	VIM 500HRpv	Kaiser Electro-Optics, Inc.	180,000 per LCD (2)	50° diag	24 oz	\$3,495
	VIM 1000HRpv	Kaiser Electro-Optics, Inc.	180,000 per LCD (4)	100° horiz, 30° vert	26 oz	\$7,995
	VFX1 HMD	Forte Technologies, Inc.	789 x 230	35.5° horiz, 26.4° vert	Unknown	\$995
	VR4 HMD	Virtual Research Systems, Inc.	742 x 230	60° diag	31 oz	\$7,900
	VR4000 HMD	Virtual Research Systems, Inc.	742 x 230	60° diag	31 oz	Contact vendor
	VRI HMD 133	Virtual Reality, Inc.	Unknown	40° horiz, 30° vert	3 lb	Contact vendor
	CrystalEyes	StereoGraphics Corp.	Not applicable	Not applicable	3.3 oz	\$595
	VR-1100	VRex, Inc.	1024 x 768	Not applicable	Not applicable	\$11,995
	VR-2000	VRex, Inc.	640 x 480	Not applicable	Not applicable	\$8,995
	BOOM 3C	Fakespace, Inc.	1280 x 1024	40° - 100° horiz	Counter-balanced	\$95,000
	Cyberface 4	Leap Systems, Inc.	480 x 640	80° horiz, 60° vert	Counter-balanced	\$15,750
	Fakespace Simulation System	Fakespace, Inc.	1280 x 1024	30° - 140° horiz	Counter-balanced	\$95,000
Autostereoscopic	PUSH	Fakespace, Inc.	1280 x 1024	45° - 110° horiz	Not applicable	\$45,000
	Virtual Window	Dimension Technologies, Inc.	680 x 480	Not applicable	Not applicable	\$17,500

There is compression of the visual field that is a function of the angle from the center: the further away from the center, the more compression.

A surprising aspect of the use of LEEP optics in HMDs is that although it was designed to be provided with a deliberately distorted image, many users make no provision for correcting such distortion. Several algorithms have been published for computing the appropriate transformation, such as that by Kalawsky (1993), and product literature from LEEP suggests correcting by making the red image about 1% larger linearly than the blue image, with the green image in between. If the eyes have full range of movement, unconstrained by the HMD, the field of view possible with the LEEP optics is 140°. With most HMDs, however, the eyes are normally constrained in their movement, with an effective field of view from 110° to 130°, depending on the amount of constraints. The focal length of the LEEP optics is 41 mm, and the entrance pupil diameter 60 mm. The LEEP optics consist of three lenses per eye.

2.2.1 Datavisor Displays

Specification	
Resolution	1280 x 1024, 1.8 arc min
Field of View	80° horiz, 20° overlap
Weight	4.2 lb
Brightness	10 fL
Contrast Ratio	100:1
Pupil Diameter	12 mm
Interpupillary Distance	56-86 mm adjustable
Image Plane Focus	Infinity to 0.5 m

Figure 2. Datavisor 10x/9ci

From n-Vision, Inc., the Datavisor 10x and 9ci devices are high resolution, wide field of view, color HMDs. They use a common display system, that is, a pair of miniature CRT displays with attached field sequential shutter devices. These shutters are electronically switchable light filters that allow only one color of light (red, green, or blue) to pass at any instant. Video is provided as sequences

of fields for red, green, and blue, with the appropriate shutter filter activated at the appropriate field. The following video formats are supported for field sequential display: 1280 x 1024, 1280 x 960, 1025 x 946, 875 x 808, 640 x 480. The Datavisor 10x unit also supports 24 bit, 1280 x 1024 60 Hz monochrome. Further details are given in Figure 2. Price information is available from the vendor.

The Datavisor VGA, as its name suggests, supports VGA video formats for field sequential display. It also differs from the Datavisor 10x and 9ci in resolution (640 x 480 pixels, 3.8 arc min) and field of view (50° diagonal). It supports standard 640 x 480 VESA and 640 x 480 field sequential video formats.

The Datavisor 80 device also differs from the above Datavisor products in resolution (1280 x 1024 pixels) and field of view (120° horizontal with 40° overlap). It supports the following video formats for field sequential display: 1280 x 1024, 1280 x 960, 1025 x 946, 875 x 808, 640 x 480. A final n-Vision, Inc. product is the Virtual Binoculars (VR-B), which is a full color, high resolution, wide field-of-view, hand held display system designed to emulate a wide range of commercial and military binoculars. Combining miniature full-color image sources and precision optical relay assemblies in a lightweight housing, it is

designed to offer the features and performance required to exploit the improved capabilities of mid- and low-range fire arms training systems as well as high end graphics training systems at a competitive cost. The Virtual Binoculars supports standard 640 x 480 VESA and 640 x 480 field sequential video formats.

2.2.2 FS5 Head-Mounted Display

Virtual Research Systems, Inc. recently announced the FS5 HMD, a display specifically designed to lower the cost of high performance VE. This HMD uses dual black and white CRTs with color shutters, and a proprietary optical design that provides a wide field of view with adjustable overlap. A standard 100% overlap allows a 55° diagonal field of view, while reducing this to 50% overlap extends the field of view to 76°. The FS5 supports full color, with field sequential RS-170 and 180 Hz RS-170 RGB. Close-cup high fidelity Sennheiser earphones are suitable for use with 3-D spatialized audio. An optional F-Scan converter drives the HMD from any mono or stereo VGA or RS-170 video source. Further details are provided in Figure 3. The price for this product is \$19,900.




Figure 3. FS5 Head-Mounted Display

2.2.3 CyberEye 100M and 100S

Specification	
Resolution	420 x 230 pixels, 3.2 arc min.
Field of View	22.5° horiz, 16.8° vert
Weight	14 oz
Brightness	Adjustable
Contrast Ratio	100:1
Interpupillary Distance	Adjustable
Image Plane Focus	~12 ft, variable

Photo courtesy of General Reality Company.

Figure 4. CyberEye 100M and 100S

General Reality Company markets two active matrix LCD CyberEye products. The CyberEye 100M is a low-cost HMD intended for extended wear, while CyberEye 100S is a stereo version of the same. Both versions support use by a single user over an extended period of time through variable focus and an adjustable interpupillary distance. Other models are available for short-duration work with multiple users or for public arcade

games; these have fixed focus and locked interpupillary adjustments. The devices support

NTSC video format. A photograph of the CyberEye display and specification details common to both the CyberEye 100M and the CyberEye 100S are given in Figure 4.

The price for the product, including audio, cables, and mounting equipment, is \$1,995 for the CyberEye 100M monoscopic system and \$2,495 for the CyberEye 100S stereoscopic system.

2.2.4 CyberMaxx CM1800

CyberMaxx CM1800 is marketed by VictorMaxx Technologies. It is a helmet-based HMD that uses a pair of color, high resolution 0.7 inch active matrix LCD displays, one for each eye. The input is standard VGA and the unit can be used with IBM compatible PCs, Macintosh computers, and with a variety of video game players. The device has adjustable interocular distance and individual eye focus adjustments. Additional specification details are given in Figure 5.




Figure 5. CyberMaxx CM1800 Display

The HMD includes a sourceless, real-time yaw, pitch, and roll head tracker and details on this component of the product are given in Figure 6. The device also includes stereo speakers with 0.1° angular resolution, 3 DOFs (yaw, pitch, and roll), and a sampling rate of 75 samples/sec. Its price is \$889.

Tracking Specification	
Sampling Rate	75 Hz
Angular Resolution	0.1° heading, 0.1° tilt
Angular Range	360° horizontal, ±45° vertical

Figure 6. CyberMaxx CM1800 Tracking

2.2.5 Dvisor Head-Mounted Display

Division, Inc., markets an active matrix LCD HMD that includes a Polhemus sensor and stereo audio. Called the Dvisor HMD, this system uses a patented depixelation technology developed by MicroSharp, that is claimed to “eliminate the pixel grid without blurring.” (The pixel grid is that matrix of pixels that is visible as individual points with low-resolution, wide field of view displays.) The optical system uses multiple aspheric lenses.

The device is described as based on an advanced, balanced ergonomic design, and designed for quick mounting and easy adjusting for any size head. Some further details are given in Figure 7. The price for this product is around \$7,000. An optional five button 3-D mouse is available with the display.




Photo courtesy of Division, Inc.

Specification
Resolution 345 x 259 pixels
Field of View 105° horiz, 41° vert
Weight 8 lb
Brightness Adjustable

Figure 7. Dvisor Head-Mounted Display

2.2.6 i-glasses!

From Virtual I/O, the “i-glasses!” is a color HMD that includes a pair of active matrix LCDs and stereo headphones. This HMD is aimed at the consumer market, with one version intended for use with video games and television, and another version that comes with a 3 DOF head tracker and an adapter that connects it to a personal computer. The device works either in a closed, immersive VE mode (with an opaque shade fitted over the outside of the headset) or a mixed-reality see-through mode. The primary difference between this HMD display and others is that the designers have chosen to sacrifice field of view for resolution. Thus, the display has only about a 30° field of view, but places the entire 640 x 480 matrix over this field. Further details are given in Figure 8. The glasses are priced at \$599 by themselves, at \$799 with tracking.

2.2.7 MRG Head-Mounted Displays

The Liquid Image Corporation (Canada) MRG 2.2, 3c, and 4 Head-Mounted Displays are all hybrid binocular, rugged displays that provide varying degrees of performance at a range of prices. These products all use full-color active matrix LCDs with an RGB delta pixel arrangement. The optics are large, 3 x 2 inches, and use a proprietary lens display. The devices include Sony stereo audio and optionally can be outfitted with Ascension or Polhemus tracking devices. The MRG 2.2 is the cheapest of these products, described by the manufacturer as the “industry workhorse” and provides a response time of 40 ms. The input video format is NTSC, with PAL optional. Further details for the MRG 2.2 are given in Figure 9. Its price is \$3,495.



GLASSES

Specification
 Resolution 640 x 480 pixels, 2.3 arc min
 Field of View 30° diagonal
 Weight 8.5 oz
 Brightness Adjustable up to 1.45 fL
 Interpupillary Distance Fixed
 Image Plane Focus Fixed at 1.8 m

Photo courtesy of Virtual I/O.

Figure 8. i-glasses!



Specification
 Resolution 240 x 720 pixels
 Field of View 84° horiz, 65° vert
 Weight 4 lbs
 Brightness 120 nt
 Contrast Ratio 40:1
 Interpupillary Distance 20 - 80 mm
 Image Plane Focus 3 ft

Photo courtesy of Liquid Image Corporation

Figure 9. MRG 2.2 Head-Mounted Display

The MRG 3c is a higher performance version display, with higher resolution and a shorter response time of 15 ms. The video signal is analog RGB and either NTSC or PAL format can be used. See Figure 10 for further details. The price for this version of the MRG is \$5,500.

The MRG 4 is similar to the MRG 3c, but intended for VE games and provides less resolution, lower contrast ratio, and a 40 ms response time. One novel feature is an optional holographic diffuser that diffuses the sharp pixel patterns into a softer image. The manufacturer claims that the device is "the world's best selling arcade HMD." Additional details for the MRG 4 are given in Figure 11. Its price is \$2,195.

2.2.8 VIM Personal Viewer

Kaiser Electro-Optics, Inc. market a low-cost HMD called the Vision Immersion Module (VIM) Personal Viewer. Intended for easy use by a number of users, this product has adjustable eyepieces, is suitable for use with eyeglasses, and has a removable head



Specification
 Resolution 768 x 556 pixels
 Field of View 84° horiz, 65° vert
 Weight 4 lb
 Brightness 120 nt
 Contrast Ratio 100:1
 Image Plane Focus 3 ft

Photo courtesy of Liquid Image Corporation

Figure 10. MRG 3c Head-Mounted Display



Specification
 Resolution 480 x 234 pixels
 Field of View 61° horiz, 46° vert
 Weight 2.5 lb
 Brightness 120 nt
 Contrast Ratio 30:1
 Interpupillary Distance 20 - 80 mm
 Image Plane Focus 768 x 556 pixels


Photo courtesy of Liquid Image Corporation

Figure 11. MRG 4 Head-Mounted Display

mount for easy sterilization. Collimating optics mean that focusing and interpupillary distance adjustments are unnecessary for each user. The VIM 1000HRpv employs four full color 0.7 in active matrix LCDs (with 180,000 pixels per LCD) with resolution limited only by these displays and not by the optics. Input is in SVGA 800 x 600 at 56-60 Hz. The VIM 500HRpv uses two LCDs, providing a reduced field-of-view, and requires a NTSC input. Both HMDs include Sennheiser stereo headphones. A photograph of the VIM Personal Viewer 1000HRpv and further details are given in Figure 12. The price for the 500HRpv version is \$2,495 and that for the 1000HRpv version is \$6,495.

2.2.9 VFX1 Head-Mounted Display System

Forte Technologies, Inc.'s VFX1 is a color HMD system that uses active matrix LCDs. It includes head tracking, audio (headphones and microphone), and a Cyberpuck that is a replacement for a mouse or joystick.



Specification
 Resolution 180,000 per LCD
 Field of View 100° horiz, 30° vert
 Weight 26 oz
 Interpupillary Distance Not applicable

Photo courtesy of Kaiser Electro-Optics, Inc.

Figure 12. VIM Personal Viewer 1000HRpv

Specification
 Resolution 789 x 230 pixels
 Field of View 35.5° horiz, 26.4° vert
 Interpupillary Distance Adjustable
 Image Plane Focus Adjustable

Figure 13. VFX1 Head-Mounted Display System

There is a large degree of software support available for the VFX1, including a CD with shareware that includes Heretic, Zephyr, Descent, Dark Forces, Magic Carpet, Quarantine, Compuserve, America Online, and 3D Ware Virtual World. It is also supported by programs such as DOOM, System Shock, SuperKarts, Flight Unlimited, and Mechwarrior2, and by vendors such as Electronic Arts, Origin, Looking Glass Technologies, and Microprose. Figure 13 provide further details about the display device. The list price is \$995.

2.2.10 VR4 Head-Mounted Display

From Virtual Research Systems, Inc., the VR4 and VR4000 are lightweight HMDs with dual 1.3 inch diagonal active matrix LCDs, as shown in Figure 14. The display system offers 10-30 mm adjustable eye relief and video input can be S-video or RGB. A position tracker is not included but can be attached. The price for the VR4 is \$7,900.

The VR4000 is similar to the VR4 but is intended for entertainment applications in which different users frequently take it on and off. It is hardened and has stereo earphones which are built in and fixed away from the head for quicker fitting. The VR4000 is intended to be sold in large quantities to OEMs and its price is negotiable.

2.2.11 VRI HMD 133

From Virtual Reality, Inc., the VRI HMD 133 is a lightweight, high resolution, color HMD that is intended for a variety of applications but particularly surgery. The unit is mounted on a headband that has ratchet adjustments on the back and top of the head. It is available in both see-through and opaque configurations, and can be easily switched between the two configurations. The video format is 1280 x 1024. Figure 15 provides additional details.



Specification
 Resolution 742 x 230 pixels
 Field of View 60° diagonal
 Weight 31 oz
 Interpupillary Distance 52 mm to 74 mm

Photo courtesy of Virtual Research Systems, Inc.

Figure 14. VR4 and VR4000 Head-Mounted Displays



Specification
 Resolution 1000 pixels per inch
 Field of View 40° horiz, 30° vert
 Eye Relief 28 mm
 Brightness 14 fL
 Weight 3 lbs
 Interpupillary Distance 62-75 mm
 Image Plane Focus Adjustable

Photo courtesy of Virtual Reality, Inc.

Figure 15. VRI HMD 133

2.2.12 CrystalEyes Shutter Glasses

CrystalEyes stereo-viewing technology, from StereoGraphics Corporation, employs electronically-switched liquid crystal lenses mounted in lightweight eyewear. The lenses shutter, either passing or blocking light, in synchrony with the left and right views of a stereoscopic image displayed sequentially on a computer or projection screen. These left and right views, called a stereo pair, are written on the screen at a high rate. When the left image of a stereo pair is displayed on screen, the switching circuit onboard the eyewear is instructed through an infrared link to cause the left lens to switch to a transmitting or clear mode. The right lens remains turned off, or in a blocking mode. Thus, the left image on screen is seen only by the left eye. During the next half cycle, the right image is displayed on screen, the right lens is switched clear and the left lens turned off. The right eye, then, sees the right image, and the left eye is blocked. The infrared link provides the synchronization signal, indicating whether a left or a right image is written on the screen, to an infrared receiver in the eyewear.

The on-screen image and the shutters switch between the left and right eye views so rapidly that the user cannot detect the shutters opening and closing. Instead, the user sees a constant image with two perspectives and the brain fuses the two images to make one 3-D image. Specification details for the CrystalEyes Shutter Glasses, and a photograph of the device, are given in Figure 16. A photograph of the glasses is shown in Figure 16, together with some specification details. The price is \$595 for each pair of shutter glasses and \$200 for the infrared emitter (which can synchronize multiple pairs of glasses).




Figure 16. CrystalEyes Shutter Glasses

2.2.13 VR-1100 and VR-2000 Stereoscopic Projection Systems

The VRex, Inc. VR-1100 stereoscopic projection system is a device for projecting 3-D stereoscopic displays on a screen and users wear passive polarized glasses to view the projected image. While the primary purpose of the device is 3-D stereoscopic display to large audiences, it also can be used as a projection display in VE systems and has been used in this way in CAVE systems, see Section 2.3.12.

The device projects via a rear-screen technique using a LCD projection active matrix display based on an 10.4 inch active matrix thin film transistor full-color LCD panel. It is designed for use with a high-end workstation, such as the Sun Sparc, but can be driven by standard NTSC/PAL and SECAM video sources. Response time is 25 ms and the display provides a resolution of 1024 x 768 pixels.

The VR-1100 projector uses a unique technology, known as spatial multiplex imaging, in which left and right eye images are combined in a checkerboard pattern. Half the pixels in the left and right eye images are thrown away, with the combining process resulting in alternative pixels that are from the left and right eye images. These pixels are polarized with a particular orientation and then viewed by ordinary passive glasses, in which one eye has a plastic sheet with a horizontal polarization and the other a sheet with vertical polarization, matching the polarization of the pixels. The advantage of the approach lies in its avoidance of the more costly and cumbersome shutter glasses, as well as reduced flicker. (The passive glasses cost about a dollar apiece.) The primary disadvantage of the technique

is the loss of resolution that results from throwing away alternate pixels. The same technology is used in 3-D LCD projector panels (used with viewgraph projectors), lower resolution projectors, and in 3-D LCD displays in laptop computers, manufactured by the same company. A photograph of the projection system is given in Figure 17. Its price is \$11,995.




Photo courtesy of VRex, Inc.

Figure 17. VR-1100 Stereoscopic Projection System

A second, related product, the VR-2000, is an integrated system that includes stereo audio. Details are given in Figure 18. Its price is \$8,995.

VRex, Inc. also markets a system known as the VR Cove, in which three VR-2000 projectors project on three screens surrounding a user.





Photo courtesy of VRex, Inc.

Figure 18. VR-2000 Stereoscopic Projection System

2.2.14 BOOM 3C

The BOOM 3C from Fakespace, Inc. is a boom-mounted visual display that consists of dual CRT displays, one for each eye, optics, and an opto-mechanical tracking system. The approach of placing a display on a boom allows the weight of the display to be counterbalanced. It provides a relatively unconstraining and comfortable interface, and one that is particularly useful with groups since it is easy to pass the display around to members of a group for viewing one at a time. The BOOM 3C uses interchangeable modules to provide a flexible display, allowing variation from 40° to 110° in the horizontal field of view. The resolution of the display is 1280 x 1024 pixels per eye, though use of standard video will

reduce this to 1280 x 960 for interlaced video and 640 x 480 non-interlaced. Video input is RGB field sequential stereo or mono. Further details on the visual display are given in Figure 19.




Figure 19. BOOM 3C Visual Display

The mechanical arm, or boom, has position sensors at six joints and the articulated support structure has 6 degrees-of-freedom (DOF) motion, to produce a global 3-D tracker for head motion. Raw analog data from the arm sensors is converted into floating point angles. Using known direct-kinematics equations, it then is possible to obtain the position and orientation of the end of the arm that supports the CRT viewer. These parameters are sent to the graphics workstation which renders the images for the two eyes so that images appropriate for the current head position can be generated. Specification details for the tracking element of the BOOM 3C are given in Figure 20.

Tracking Specification	
Sampling Rate	>70 Hz
Latency	200 ms
Accuracy	0.16 in
Resolution	0.1°
Working Volume	6 ft diameter circle, 2.5 ft vertical
Field of View	140°

Figure 20. BOOM 3C Tracking

This device costs \$95,000, including both the tracking and visual display elements. Custom versions of the BOOM 3C are available. MedView, for example, is a custom version of the BOOM 3C that Fakespace, Inc. developed for medical purposes.

2.2.15 Cyberface 4

The Cyberface 4, from Leep Systems, Inc., is an active matrix LCD boom-mounted display, and is an extension of a previous product by LEEP systems known as the Cyberface 3, which is available now only on a custom-order basis. The display uses high resolution 12-bit VGA (RGB) active-matrix LCDs to provide relatively high resolution over a wide field of view. The pixel pitch is 6.2 arc minutes, while the pixel structure is vertical stripe. Further details are given in Figure 21. (Note that the figures given for brightness and contrast are preliminary.) The price is \$15,750.


Another product, the Cyberface 5, is expected to be available in 1996. This will be a more sophisticated product using two LCD displays for each eye, one that has high resolution for the fovea and another that provides low resolution with a very wide field of view. The images from the two displays are optically added together. The display weight is expected to be 26 ounces, and the estimated price is \$45,000.

Specification	
Resolution	480 x 640 pixels, 6.2 arc min.
Field of View	80° horiz, 60° vert
Weight	Counter-balanced
Brightness	~3 fL
Contrast Ratio	≥80:1
Pupil Diameter	70 mm
Interpupillary Distance	64 mm
Image Plane Focus	42 mm effective focal length

Figure 21. Cyberface 4

2.2.16 Fakespace Simulation System

Another Fakespace, Inc. product, the Fakespace Simulation System (FS2), is a counter-balanced boom-mounted immersive display that uses dual CRT displays. The electronics, optics, and mechanics of the device are similar to the Fakespace BOOM 3C, but the FS2 is designed for immersive rather than pass-around viewing. Like the BOOM 3C, this device has 6 DOF movement and optics with user interchangeable modules that allow a horizontal field of view ranging from 30° to 140°. Video input is RGB field sequential color and the video resolution is 1280 x 1024 interlaced or 640 x 480 non-interlaced. A photograph of the device is given in Figure 22, together with some specification details. The product is priced at \$95,000.



Specification	
Resolution	1280 x 1024 pixels
Field of View	30 - 140° horiz
Weight	counter-balanced

Photo courtesy of Fakespace, Inc.

Figure 22. Fakespace Simulation System

2.2.17 PUSH

The Fakespace, Inc. Personal Use Stereoscopic Haptic (PUSH) display provides an immersive display for individual use. The device uses one CRT display for each eye and an opto-mechanical tracking system. Based on the BOOM 3C technology, also manufactured by Fakespace, Inc., PUSH is built on an 18 inch, 3 DOF desktop support structure, provid-

ing full 6 DOF control. PUSH features an easily mastered, intuitive method of controlling movement in the VE by simply pushing or rotating the device in the desired direction. A photograph and specification details are given in Figure 23. The PUSH is designed for OEM application developers and costs \$45,000.




Figure 23. PUSH


2.2.18 Virtual Window

Dimension Technologies, Inc.'s Virtual Window is an LCD backlit display using an autostereoscopic technique. The device is based on the parallax barrier approach, in which a mechanical or optical arrangement is used so that a given pixel is only seen by one eye. The Dimension Technologies approach is unusual in that instead of using lenses or vertical bars in front of the display surface to prevent more than one eye from seeing a pixel, it uses a unique backlighting arrangement that accomplishes the same goal. Rather than having an even backlit illumination, the device uses a backlight consisting of a large number of narrow vertical lines equally spaced, with one vertical line for every two columns of pixels in the display in front. The effect of this is that a given pixel can be seen only by one eye, and not the other, like parallax barrier systems, but without the reduction in brightness that results from cutting off part of the light.

This display is claimed to be the only autostereoscopic display device currently on the market. The device is full color, with 6 million colors, and advertised to have brightness 1.6 times that of a standard CRT. It can be used in either a 3-D mode or a 2-D mode. The display size is a little over 11 inches diagonal. Further details are given in Figure 24. The price for Virtual Window is \$17,500.

2.3 Current Research and Development

The following research laboratories and small companies are performing research related to visual displays for VE systems. These can generally be divided into two different categories: (1) the development of new types of 3-D displays, and (2) research into



Specification
 Resolution 680 x 480 pixels
 Weight Not applicable
 Brightness 72 fL
 Contrast Ratio 100:1

Photo courtesy of Dimension Technologies, Inc.

Figure 24. Virtual Window

how to use these displays, including investigation of human factors issues. By far the most active research activity at the present time is the attempt to develop autostereoscopic displays. This is because the basic technology for constructing HMDs and shutter glasses is already well known and available as commercial products. The autostereoscopic systems described below typically have multiple applications, including both VE systems and telepresence systems, and more ordinary uses such as 3-D computer displays or television.

The extent to which HMDs and shutter glasses are likely to be replaced by autostereoscopic systems is unclear, and depends upon the particular application and how much of a problem head-mounted gear is for that application, as well as the performance and cost of those autostereoscopic systems that are developed into commercial products. However, it is clear that autostereoscopic systems can be used not only for non-immersive applications, but also immersive applications by using large-screen projection displays, similar to those described in the CAVE project, discussed in Section 2.3.12 below, that is being undertaken by the University of Illinois at Chicago.

2.3.1 ATR Communications System Research Laboratories, Japan

The ATR Communications Systems Research Laboratories are developing virtual space teleconferencing systems in which participants experience a sense of telepresence—in this context, the sense that they are physically in the same environment as other participants they are communicating with, even though they are geographically separate. The researchers are led by Katsuyuki Omura.

One specific objective toward their long-range goal is the development of autostereoscopic display systems. This type of display was chosen in preference to holography and volumetric approaches because of the requirement for full color, a large display size, and imaging in real time. The researchers have chosen an approach using lenticular screens in which a display screen is overlaid with a sheet of tiny lenses so that a given pixel can be seen by only one eye. The primary problem with the basic lenticular approach

is that viewers are not free to move their heads laterally beyond a trivial amount. Movement beyond the immediate area results in either no view (because the viewer is seeing the dead space between pixels) or a reversed 3-D effect if the viewer moves further laterally. The ATR group is resolving this problem by tracking the viewer's eye position so as to change the display appropriately when his eye moves into another area. Previous work demonstrated the feasibility of this for a single viewer and a high-density LCD projector.

Recent work has focused on extending the approach to multiple viewers. In this more recent work, the position of each of several viewers is tracked using magnetic sensor or other techniques, and the image projection modified according to the movement. The basic approach is that for every viewer, there is a projector that projects both a left and right image on a screen. The viewers are on one side of a large (100 in) screen, while the projectors are on the other, creating a rear-projection image on the screen. Each projector is mounted on guide rails that allow it to move (driven by motors) a maximum of 0.9 m laterally and from 1.2 to 2.5 m in the front and rear directions: the position of each projector typically mirrors the position of the viewer it projects images for. If the viewer moves to the left (from the viewpoint of the viewer), the projector will also move to the left (from this same viewpoint). If the viewer moves towards the screen, the projector also moves towards the screen. This mechanical arrangement does substantially restrict the mobility of viewers. In the prototype system, only two viewers can be accommodated, because the size of the projectors allows only two projectors, although the researchers envision more viewers in an eventual product. Each projector uses a specially developed projector lens with a wide 65 mm exit pupil, and uses a CRT rather than an LCD because of the CRT superiority in wide angle projection. The projection screen consists of two layers of lenticular lenses, with a diffusion layer between the layers. Tests of the system confirmed that two viewers could in fact see different stereoscopic images. The researchers expect future work to involve developing a system for determining viewer eye positions by the use of video cameras.

As part of developing systems that have high fidelity telepresence, the researchers are studying certain errors in the perception of stereoscopically presented objects. This is motivated not only by the goal of allowing participants to perceive themselves as being present at the same location as other conference participants, but also by the goal of being able to manipulate objects in that environment with the same dexterity that they could if they were physically present. The approach taken is to use a large (709 in) stereoscopic projection system, using the autostereoscopic techniques discussed above, rather than HMDs, which the researchers consider too cumbersome. The focus of the study concerns the visual perception of objects in this space, in which both virtual and real objects exist, and in particular the need for accurate perception of the location of objects to allow a human to reach out with his hand and "grab" a virtual object. This requires the perception of an object location to be consistent with the representation held by the computer controlling the VE, and also requires the perception of the object to be stable when the user's hand approaches it. This can be difficult because in environments in which virtual and real

objects are mixed, a variety of phenomena occur, including the following: (1) conflict between accommodation of a viewer's eye at the location of the real or virtual object, as opposed to accommodation at the location of the screen; (2) the tendency for real objects to become transparent when they pass behind a virtual object; (3) the differences in hue, brightness, contrast, and similar characteristics between real and virtual objects; and (4) errors in tracking the user's head and hand positions.

To date two main issues have been focused on. The first concerns the depth error resulting from a mismatch between the assumed inter-pupil distance (IPD) and the actual IPD of a viewer. It is well known that a mismatch between the assumed ("setup") and real IPDs of an observer can result in significant depth error. What has not been previously investigated is the extent to which this mismatch and resulting depth error can be caused by convergence, in which the individual eyes rotate to converge on an object at a given distance. This rotation results in changing of the IPD. The researchers measure the error in depth perception as a function of the position of the (virtual) distance between the displayed object and the screen it is displayed on. A mathematical model predicted that if in fact there was error due to a change in IPD resulting from convergence, there would be a systematic, linearly increasing error of a certain magnitude. In fact, the actual error was more than twice that of the predicted value, suggesting that while changes in the IPD resulting from convergence may be part of the source of error, there are other sources of error. The magnitude of these errors is viewed as unimportant for large-screen applications such as teleconferencing, but potentially quite significant in the case of see-through HMDs.

The second issues that has been examined is the effect of fuzziness of a displayed image on depth perception. In an experiment, the researchers had subjects make judgments about the depth of sharp and blurry images. Results showed that all subjects perceived the blurry image as farther away than the sharp image, though there were large individual differences among subjects, with the effect very small for some subjects.

Current and future research is focusing on the following: (1) determining the other source of depth error (other than convergence); (2) determining what factors are involved in depth error that may interact with blur; and (3) investigating the effect of fuzziness in mixed virtual and real environments.

2.3.2 British Aerospace plc, United Kingdom

Researchers at the Sowerby Research Centre, British Aerospace, plc, are investigating the extent to which viewers can adapt to the unusual accommodation that is typically necessary when using HMDs and heads-up displays. Accommodation is important, in that improper accommodation can result in blurred objects and a failure to detect objects. A variety of factors can influence accommodation. In darkness, accommodation moves to a resting position. If a stimulus is blurred, a reflex drives a change in accommodation in an attempt to resolve the blur. Accommodation also tends to change if there is a

change in vergence, the lateral movement of each eye that causes the two eyes to track together. The cognitive knowledge that an object exists that is close to the eye may cause accommodation, and specific conscious mental effort can result in a lapse of accommodation. The experiments being performed by the British Aerospace researchers are intended to resolve how these factors work together when HMDs and HUDs are used.

In an initial series of experiments, the researchers measured resting accommodation in darkness, finding that subjects focused at about 1.4 Diopters, or about 0.7 m. When an optical combiner was placed directly in front of the subject's eyes, the accommodation did not change, even though subjects were aware of the combiner being in front of their eyes. Then an array of hash symbols that formed a sharp, high-contrast pattern was projected as a virtual image on the combiner, collimated such that subjects should accommodate at infinity. This particular image was chosen since it should serve as a high-quality stimulus that triggered an accommodation reflex. Only three of the subjects were able to maintain accommodation at or near infinity, with the other five accommodating at various levels, including as short as 1 m. When the virtual image projected on the combiner was changed to a word, all subjects showed an accommodation substantially closer to the subject than with the previous image, even in the cases where the earlier accommodation was at or near infinity. These experiments suggest that subjects are substantially misaccommodating when virtual imagery is present, particularly when they are mentally processing information in the imagery.

In another experiment, subjects were tested for their ability to accommodate to virtual imagery when it was superimposed on the real world. In this experiment, subjects looked out an open window to see a brick wall and bushes about 28 m away. Three conditions were run. In one, subjects were asked merely to view the scene and to keep the wall and a light fitting on the wall in focus. In the second condition, an array of hashes was superimposed on the outside world with a beam splitter, with the virtual image collimated so that it appeared to be at the same optical distance as the wall. In the third condition, reversed words were presented in the array of hash marks, and subjects were asked to read the words aloud. In the first two conditions, most subjects could maintain accommodation at or near infinity. However, when subjects were required to read the reversed words aloud, thereby mentally processing the information in the virtual image, every subject showed a lapse of accommodation inward, in most cases, one of quite substantial magnitude. There was little difference between accommodation by subjects reading words aloud when they saw a mixed real plus virtual scene than when they saw the virtual simulation. The implications of this experiment are substantial: it suggests that subjects using a HMD will not accommodate to infinity, if intended by designers, even if they have a mixed scene and real-world stimuli to focus on. Rather, their accommodation will lapse, resulting in blurred objects and the potential of failure to detect objects. Such misaccommodation can also result in misperception of the size and distance of an object.

In a follow-up experiment, subjects were provided with an information processing task presented either visually or aurally while viewing a simulated scene, and the shift in

accommodation measured. Subjects shifted their accommodation in both cases, but less so when information was presented aurally. This suggests that it might be better to provide a mixture of information to persons viewing HMDs, with a substantial part of it aural, to reduce misaccommodation. However, it is possible that the experimental results could be due to the visual task being more difficult than the auditory analogue.

Other work at British Aerospace includes a study of why blurred images appear sharper when in motion, and studies of the conditions under which the movement of an image on a display is perceived as smooth movement by a viewer rather than jerky movements or multiple images. Still other work includes investigation of voluntary head movements during visual tracking and the resulting slippage of a helmet, and the use of eye pointing as an input media.

2.3.3 BT Laboratories, United Kingdom

Researchers at BT (formerly British Telecom) Laboratories are looking at how useful 3-D might be in video-teleconferencing systems. They see using glasses in teleconferencing applications as very undesirable, "which dramatically reduce eye contact" and which can "make the wearer look doubtful or sinister. Since eye contact with the person at the remote location is one of the key advantages offered by video-telephony and video-conferencing, spectacles based 3-D imaging approaches are not appropriate for these applications" (Jewell et al, 1995). The work of these researchers with 3-D displays also is applicable to telepresence systems and VE applications in which it is undesirable to have the user wear a HMD or shutter glasses.

The researchers are using a system in which they have positioned a lenticular sheet in front of an LCD display, where the sheet is viewed at a distance of 600 mm. (The lenticular display was chosen in preference to a parallax barrier because of the higher luminance throughput.) In such a system, the lenticular sheet consists of a set of columns of tiny lenses, with each column mapping to a pair of columns of pixels, one column of pixels displaying the image for the right eye and the other displaying to the left eye.

Early work looked at necessary properties of the LCD display used with the lenticular sheet. The geometry of the early display resulted in dark areas between pixels that were perceived by the viewer as dark stripes, particularly when the viewer turned his head. An LCD with a very limited color palette (3 bits per color) showed good quality 3-D but rough transitions between shades were particularly distracting in the case of skin tone. A newer display with 200,000 colors has provided good 3-D and much better color, and is currently in use. Preliminary tests have also looked at the transmission of display data over communication lines. Two separate channels were used with independent compression hardware and, though there is a potential problem if the two transmitted views become unsynchronized, overall results were good.

More recent work has concentrated on the development of a head tracking system to allow viewers to move around and still see an appropriate 3-D effect. The general prob-

lem with lenticular systems is that when a viewer moves laterally, movement beyond the immediate area results in either no view (because the viewer sees the dead space between pixels) or a reversed 3-D effect with further lateral movement. The BT researchers used a commercially available infrared head tracking system that tracked the lateral position of the viewer's head. They compared moving the lenticular sheet with respect to the LCD display, to correspond to head movements, against rotating the entire assembly of LCD display and lenticular sheet. Of these two approaches, the second proved superior. Since viewers seated in front of a video telephone rarely moved more than 250 mm on either side of a central resting position, and these movements were usually slower than 1.35 m/sec, the necessary rotations can be supplied by a standard stepper motor.

In the long term, these researchers expect that using video cameras with image processing hardware and software that is capable of locating and tracking individuals, together with displays with less dead space between pixels, will generally solve the problem of limited viewing areas for lenticular systems. They would also like to see greater resolution for displays, and image processing systems that can interpolate between source camera views to produce a greater number of intermediate images for viewers to see as they move their head. The goal of present work is extending their basic approach to a higher resolution display. However, they see the most practical future approaches to 3-D video telephony generally as not involving head tracking, but as using high enough resolution displays, together with some additional bandwidth to provide additional perspective views. The combination of additional resolution (which allows more movement laterally by viewers) and additional perspective views would eliminate the need for head tracking.

2.3.4 Canon, Inc., Japan

Researchers at Canon, Inc. are developing techniques for creating a large number of different viewpoint stereoscopic images from a smaller number of 3-D images, or creating 3-D images from 2-D ones by interpolation. The ability to present different perspectives of an image is needed for any type of display that is viewpoint dependent, for example, binocular displays and autostereoscopic displays using lenticular and parallax barrier approaches.

The Canon researchers have developed algorithms for creating interpolations of scenes from a unique perspective between two given perspectives. These interpolations have been experimentally tested by presenting them to viewers using a CRT monitor and shutter glasses. The interpolation method starts by constructing a data structure known as an epipolar-plane image (EPI), in which the separate images from a line of cameras are matched up together into a volume such that the position of the camera is the third dimension. If the cameras are aligned on a straight line, a point in 3-D space is seen on the EPI as a straight line, known as a trace line. All of the possible trace lines that pass a point are identified by searching for trace lines that have similar color values and a slope within a certain expected range. Every pixel that is not already part of a trace line is processed to locate a trace line it is part of. Then the view from a virtual camera is created by determin-

ing the value a trace line would have if it passed through the space of the virtual camera position. This step imposes a number of difficulties. One such difficulty occurs when two or more trace lines can intersect at a given point on an EPI, posing the problem of which line to use for an interpolation. In such a case, the line with the least slope is selected to use for interpolation. Another problem is posed by a background region with uniform color that can produce many trace lines and prevent selection of a true trace line, resulting in the background hiding objects in front of it. The algorithms developed can handle these and other problems.

The algorithms have been applied to a number of real scenes, including complex scenes that resulted in considerable motion parallax. In addition to the algorithm for creating a virtual camera position by interpolating between camera positions along a straight line, the researchers have developed an algorithm for reconstructing back-and-forth multi-viewpoint images from a set of right-and-left view point images. These are not simply images that one might see by zooming in and out, but perspectives as might be seen by actual physical backward and forward movement.

Future work is expected to focus in two areas: (1) speeding up the interpolation process and making it more robust when operating in a range of different conditions; and (2) extending the algorithms to create VEs that consist of both real and computer-generated imagery.

2.3.5 Dimension Technologies, Inc.

Researchers at Dimension Technologies, Inc., are developing, with funding from the National Aeronautics and Space Administration (NASA), a prototype autostereoscopic display that produces multiple perspective views with full resolution. The overall goal of the work is to develop a device that is not only autostereoscopic, but has a form known as "look around" in which viewers see different perspectives as a result of moving their heads.

Initial work developed proprietary technology that uses the parallax barrier technique for autostereoscopic display. This technique involves placing a barrier between the viewer's eyes and the screen, such as a sheet of opaque material with narrow vertical slits. Such a barrier prevents both eyes from seeing the same column of pixels in the image, and allows presentation of different stereoscopic images to the left and right eyes. Dimension Technologies, Inc. has developed a form of this technology that doesn't actually use a barrier, but narrow vertical light lines that are positioned behind a LCD. The light backlighting the LCD consists of a substantial number of such vertical lines, spaced at equal horizontal intervals. The geometry of the situation is such that there is a vertical line of light for each two columns of pixels, allowing each eye to see only one column. The effect is similar to that of an opaque barrier with a vertical slit, but allows greater brightness and the ability to turn the vertical lines on and off as required.

This basic technology is used in the Virtual Window commercial product

described in Section 2.2.18. Its primary limitation is that the reliance on the geometry of parallax creates specific viewing regions for the left and right eyes. If the user moves his head laterally a sufficient distance, he will leave the correct viewing regions and see either no 3-D effect (if at the edges of the regions for a particular eye) or, perhaps worse, an inverted 3-D effect (if the left eye is in a region intended for a right eye, and vice-versa). Therefore, the viewer must remain stationary or, more typically, a system must be provided for tracking the lateral position of the head and modifying the display accordingly. While this solution can work well for a single viewer, the fact that the display can only be adjusted to movement of a single viewer means that it works poorly for multiple viewers.

In the ongoing research, a number of enhancements are being added to the system. First, a very fast LCD was developed allowing individual pixels to be addressed and pixel values modified at a rate several times faster than current off-the-shelf LCD products. A prototype LCD, built by the David Sarnoff Research Center, has an approximate 180 Hz address rate and pixel response times of 0.5 ms off and 3.5 ms on. Second, an illumination system was developed that consisted of 24 fluorescent lamps and an opaque sheet with narrow vertical lines to allow transmission of light to the LCD. Third, a computer control system is used for the lamps that allows each lamp to be turned on and off selectively. Fourth, a sheet of lenticular lenses, consisting of 265 parallel, vertical cylindrical lenslets molded of plastic, and bonded to glass, is placed behind the LCD displaying the image. Fifth, a memory system provides for storing images received at 60 frames per second, so that they can be displayed on the LCD at 180 frames per second.

The resulting prototype display system has so far been able to produce six perspective views at 800 x 400 pixel resolution, with a 32-level grey scale, where each perspective view is visible within a viewing zone 6 cm wide. A viewer thus sees a single view and, if he moves laterally, will see the same perspective until moving into the next viewing zone, in which case a different perspective view is seen. Thus, 3-D images can be seen across a 36 cm viewing area. Brightness is 23 fL and a 15:1 contrast ratio was measured. Ghosting, a form of crosstalk where images intended for zones other than the current one are faintly visible, is relatively low, about 5% as bright as the proper image. The change in perspective as the viewer moves from one zone to another also is relatively low, implying a smooth "look around." The measured critical flicker frequency is about 36 Hz high, and may be reduced with better equalization of the brightness of the different lamps.

Most of the technologies used in the prototype for electronics, lighting, and optics are proven variations on standard technologies, and Dimension Technologies, Inc. believes that the primary barrier to commercializing the prototype is the availability of higher resolution, higher speed, color LCDs that can be mass produced. They are engaged in discussions with manufacturers towards this end. They also expect that very small LCDs that have recently appeared on the market for use in projection systems can be adapted to autostereoscopic displays using their approach.

2.3.6 Dimensional Media Associates

Dimensional Media Associates is developing autostereoscopic 3-D volumetric displays that could be used in VE systems. Images acquired from objects, CRTs, LCDs, or other media, are projected as full color, solid objects floating in midair. This technology is being sold as a product, the High Definition Volumetric Display (HDVD), for specialized applications such as displays in retail stores where the effect is attention-getting.




Figure 25. HDVD Overview

The HDVD technology is based on a technique described in a US patent (Summer & Katz, 1994) that uses a pair of large concave mirrors packaged with an image source (an object itself or a 2-D video screen), such that the image source is reflected by first one mirror, then the other, so that the object appears to be floating in space (see Figure 25). The image source in the current HDVD system is provided by 2-D display signals, such as those available from video sources or computer

simulations. It is, therefore, only capable of projecting monoscopic images. However, the approach is capable of projecting a 3-D image if a volumetric display is used as the image source. Dimensional Media Associates is adapting an acoustic-optical scanner, a volumetric display that uses audio frequency to change the index of diffraction of a lens to diffract the light beam, for use with the concave mirror HDVD technology. The company describes the technology as having potential for application in information kiosks, medical diagnostics, air traffic control, video games, large-scale theme park rides, and data visualization. The company has a Small Business Innovation Research contract from ARPA to develop applications of their display technology to the simulation of surgery.

2.3.7 IBM Thomas J. Watson Research Center and Georgia Institute of Technology

James Lipscomb of the IBM T. J. Watson Research Center, in collaboration with Wayne Wooten of the Graphics, Visualization, and Usability Center at the Georgia Institute of Technology in Atlanta, is researching image processing techniques for reducing crosstalk in shutter glasses. In display systems using shutter glasses, crosstalk occurs when the image presented to one eye is unintentionally seen by the other resulting from leakage through afterglow of the phosphors on the CRT and leakage through the LCD shutter. The first case arises because while the entire LCD shutter turns on or off at one time, the same is not true of the image on the CRT viewed through the shutter. The luminance is turned on by the trace of a scanning electron beam that moves horizontally from left to right and then retraces additional horizontal lines until it reaches the bottom of the screen. It is turned off, effectively, when the light from phosphors on the screen decay sufficiently, in the same time sequence as the tracing beam. In the second case, leakage through the LCD shutter happens when a shutter begins to go black during the vertical blanking interval of the video signal but, when the next field begins to be traced by the

electron beam, the shutter has not completely gone black and is still settling down. These characteristics mean that crosstalk varies considerably as a function of the vertical position on the screen. As measured by Lipscomb and Wooten (1994), the leakage through the shutter at the top of the screen (about 7% leakage), continues about 15% of the way down the screen (4.5% leakage). From then on, crosstalk actually increases due to phosphor persistence showing an afterglow from the electron beam tracing, and this crosstalk rises exponentially, reaching a maximum at the bottom of the screen.

Lipscomb and Wooten have developed an algorithm for minimizing crosstalk arising from these causes. First, the image is processed so that the brightness of each pixel ranges from 0.3 to 1, rather than 0 (black) to 1 (white). By reducing the darkest normal intensity level to a dark grey, a certain level of crosstalk can be eliminated since it will not show up against the lighter background. (This strategy incurs the penalties of reducing overall contrast and making images light pastel colors). Crosstalk is further reduced by predicting the amount expected for each eye and subtracting that level from the image. Since crosstalk varies greatly by color, because phosphors that produce some colors have greater persistence than phosphors that produce other colors, this subtraction is color specific. For example, blue and green phosphors have zinc sulphide, resulting in a longer persistence than red, which does not. Consequently, image subtraction is done at one level of magnitude for green and blue, but at a much lower level for red. The algorithm also caters for crosstalk differences based on the vertical dimension of the screen. In this case, the screen is divided into sixteen horizontal bands, with anti-crosstalk measures greatest in the top two and bottom four bands, and weakest in other bands.

2.3.8 Infinity Multimedia

Infinity Multimedia is developing an autostereoscopic display system based on research conducted at the Computation Laboratory and Engineering Department of the University of Cambridge in England. Instead of the more common parallax barrier or lenticular approaches, this display makes use of a patented beam steering system in which images are made visible to only one eye at a time.

The beam steering system allows viewer head movement by using multiple images that are presented via time-division multiplexing. Thus, rather than displaying 60 fields per second, the system displays six different perspective images, using a total of 360 fields per second. This significantly broadens the viewing area that viewers can move around in, allowing the viewer to move forward and backward with respect to the screen as well as laterally. The basic approach uses a CRT, a pair of lenses, and a ferro-electric LCD. The CRT display produces an image, with the light produced by the display passing through the first lens and then through a narrow horizontal slit displayed on the LCD. The light then continues through a second lens to an eye of the viewer. The slit is in the focal plane of the second lens and this arrangement results in all of the light that forms the image on the CRT display being transmitted from the lens in a single direction, and thus viewable only from a single direction. In particular, the image is viewable by one eye but not the

other. Changing the position of the slit (by an appropriate display on the LCD) changes the direction (and thus the eye) from which the image can be viewed. A stereoscopic image can thus be presented by displaying the appropriate image in synchronization with the appropriate slit position for each eye.

Infinity Multimedia presently has a 10 inch diagonal CRT-based proof-of-concept prototype built out of standard components. By spring 1996, the developers expect to have completed an engineering prototype, with a 25 inch diagonal display, using specially designed components. Full-scale production of a commercial product is expected by the end of 1996. Both the engineering prototype and production system are expected to replace the CRT display with LCDs and rear-projection methods. These systems require special high-speed LCDs, and one of their corporate partners, Litton Industries, is contributing the necessary technology, including both lens design and very fast LCDs based on Cadmium Selenide active matrix transistors.

2.3.9 NASA Ames Research Center

Researchers at the NASA Ames Research Center have continued the human factors research in VEs that was initiated in the 1980s with their pioneering efforts in creating the first HMDs. The recent work has had a variety of goals.

One goal of the research is that of developing techniques for calibrating displays for VE systems. Calibration is desired for two reasons. First, it is desirable to ensure that the viewing optics have a field of view that matches the field of view expected by the graphic display system. Second, the optics of HMDs, particularly those with low-cost lenses, introduce distortion. While algorithms exist that can predict the amount of distortion, these algorithms assume an idealized viewing situation and their accuracy in an actual system is unclear. In see-through displays, calibration is not a problem because computer-generated imagery can be superposed on actual physical objects, and the imagery aligned to the physical objects by vernier adjustments. However, in closed systems, some alternative technique is required that can calibrate the imagery and verify that the calibration and registration is correct. Work by Stephen Ellis and Kenneth Nemire has used psychophysical techniques that involve subjective judgement of visual direction. Thus, in an experiment, subjects were asked to point in the direction of either a real object or, in a different condition of the experiment, a virtual object, with both objects appearing as fence posts. Subjects were highly accurate in pointing to the physical object but much less accurate in pointing to the virtual object, a difference the researchers attributed to an incorrect scale factor for the viewing angle in the virtual condition. After modifying the scale factor, accuracy improved considerably. The researchers also attempted to align the virtual and physical environments together with what the subjects perceived as being straight ahead, but found that this did not improve the rotational errors, and concluded that some other factor must be responsible for this error. They are engaged in further experiments in an attempt to isolate the cause of this error.

Another line of research at NASA Ames is a test of whether the addition of a third dimension in a head-slaved telepresence situation (in which an operator is provided with a HMD and a remote camera that moves in a manner slaved to the operator's head movements) can enhance the awareness of a spatial situation. Traditional telepresence systems use two degrees of freedom: pan (azimuth, yaw) and tilt (elevation, pitch). Recently, there has been a trend toward adding "roll," so as to mimic the full capabilities of the human head. The issue here is whether the addition of roll is worth the additional cost and complexity required for its use. Bernard Adelstein and Stephen Ellis carried out experiments that compared the performance of subjects both with and without roll capability. They found that the ability of subjects to accurately determine the azimuth or elevation of an objection was not improved by the roll capability, but that the ability to determine the orientation of an object was improved by a roll capability.

In another line of research, the effect of the display of an object as pitched—that is, rotated up or down from the horizontal plane—on an observer's perception of gravity referenced eye level was investigated. Previous work had shown that when the display of a virtual box was pitched up or down, subjects' judgments of eye level were biased in the direction that the box was pitched. Recent work has shown that the magnitude of this effect depends upon the structure of the object being displayed. In addition, the effect was not as strong with a virtual box as with a real box. The researchers also found that longitudinal structure biased the perception of eye level more than did traverse structure. While the perceptual effect was not as strong with a virtual box as with a real box, only minor additions to the display of the virtual box were necessary to obtain the effect comparable to that of a real box. They found that observers adapted with experience by tending to increase their perceptual bias. Future experimentation may better reveal how observers adapt to the VE.

Other research at NASA Ames focuses on perceptual phenomena that can degrade perceptual performance. In particular, work is ongoing on the source of errors in the perceived distance to virtual objects. Here, the researchers found that superposing a virtual object on a physical backdrop changes its position as judged by an observer. Specifically, if a physical surface is introduced at the depth of the stereoscopic virtual image of an observer, the virtual object is judged to be closer to the observer.

2.3.10 Purdue University

Researchers at the School of Electrical Engineering, Department of Psychological Sciences, and the Biomedical Engineering Center at Purdue University are engaged in experimental studies to determine the extent to which viewing images in stereo can support better visual task performance than viewing images monoscopically. In particular, they are investigating whether the presentation of X-ray images as fused stereo pairs can provide radiologists with depth information that is similar to that obtained with techniques such as computed tomography, but with considerably less radiation dosage and less cost.

The researchers have experimentally tested the ability of subjects to detect a potential tissue abnormality, that may signify a small early cancer, from breast X-rays. In the two experimental conditions, images were displayed on a CRT display while viewers wore LCD shutter glasses. In the stereoscopic condition, the views for the left and right eyes were displayed in alternate video fields. In the monoscopic condition, images were presented side-by-side, so that depth was not perceived. In half of the trials, the task was to decide whether an object of "higher density" was present, while in the other half the subjects were asked to decide whether a particular target arrangement of objects was present. The results of this experiment showed that the subject performance in detecting high density objects was comparable under stereoscopic and monoscopic conditions, but the stereoscopic presentation did increase subject performance in detecting a specific arrangement of objects.

2.3.11 Terumo Corporation, Japan

The Terumo R&D Center, Terumo Corporation, in collaboration with the Department of Radiological Technology, Nagoya University of Medical Technology, Japan, is developing an autostereoscopic display using LCDs. The display system uses an unusual method of head tracking and providing for a parallax barrier. While many variations have been reported, the basic idea is that of having an infrared-sensitive television camera capture the image of viewers illuminated by an infrared light. This image then is displayed on a black-and-white display screen that serves as a backlight to a color LCD that displays the actual image to be viewed. This scheme serves the same purpose as a head tracking system and a parallax barrier system—allowing a viewer to see a particular perspective only from a certain location.

In the first version of the system, a single monochrome display is used with a single TV camera but a pair of infrared lamps, one illuminating viewers from the left side, the other from the right. The stereoscopic signal is time sequential, with alternate fields viewed by the left and right eyes. This is accomplished not by shutter glasses, as is conventional, but by turning on each infrared light according to whether the field is for the left or right eye. If it is for the left eye, the lamp illuminating the left side is turned on. This results in the image of a half-face on the monochrome display, which backlights a color LCD with the actual image to be displayed.

An alternative system is a time-parallel system. This has a pair of displays, each consisting of a black-and-white LCD, a large format convex lens, and a color LCD, with the lens set up so that the black-and-white LCD backlights the color LCD, upon which is displayed the image to be seen. Each of the pair of displays is arranged at right angles to each other, with a half-silvered mirror at a 45° angle such that the two images are combined as seen by the viewer. The viewer is illuminated by a pair of infrared lamps, one from the left and one from the right, as before, but in this case each lamp has an infrared filter in front of it, either 830 - 870 nm or 930 - 970 nm. A pair of television cameras are used, with one of the indicated infrared filters in front of each camera, and the output of

each presented to the black-and-white LCD of each display. The image displayed on the black-and-white LCD is a half face image of each viewer. (In the case of the right eye, the right half face).

Still another system uses a similar approach but allows multiple viewers, is thin, and can be mounted on a wall. It is expected that some version of the Terumo Corporation system soon will be introduced in the United States as a commercial product.

2.3.12 University of Illinois at Chicago

For several years, researchers in the Electronic Visualization Laboratory at the University of Illinois at Chicago have been developing a VE and scientific visualization environment in which a person, or group, is surrounded by screens that provide visual displays. The environment is known as the CAVE, a recursive acronym that stands for CAVE Automatic Virtual Environment. Displays are projected on screens positioned at the front, two sides, and floor of a 10 x 10 x 10 ft room. (A sphere would actually be better, so the VE would be seamless, but this is computationally very expensive and beyond the capability of present graphics hardware.) Rear projection is used for the three walls and down-projection for the floor. Each screen uses a separate Silicon Graphics high-end workstation to create the graphics, providing a color display with a resolution of 1280 x 512 pixels. The visual display is at 120 Hz, with alternating fields for different eyes and users wear StereoGraphics CrystalEyes shutter glasses. Multiple speakers provide 3-D sound, and users wear electromagnetic sensors that track head and hand movements.

There are a number to advantages of the CAVE over HMDs, including the ability of users to easily see others in the same room, some physiological vision effects such as the ability to see objects appropriately in focus or out of focus, and significantly less occurrence of motion sickness: the researchers report that only two of the 9000 people who have visited the CAVE complained of nausea. The CAVE display also is claimed to provide a more accurate display than HMDs, which have optics that create geometrical distortion. In designing the CAVE the researchers had to confront a number of problems, including the difficulty in displaying green stereoscopically by projection (caused by especially long persistence of phosphors in projection equipment, and solved by a specialized tube), minimizing user shadows when projecting downward onto the floor (shadows cannot be eliminated unless the floor has projection from below, but projection from the top offset to the front minimizes shadows), the use of a shared memory arrangement among workstations to synchronize frames on different screens so as to avoid a problem in which "images in the corners crease and start to look sucked in like sofa cushions."

The researchers analyzed the HMD, monitor, and CAVE situations with respect to the results of tracking errors that are errors of displacement or errors of rotation. Neither monitors nor the CAVE are sensitive to rotation error, since the image display plane does not move with the position and angle of the viewer. In the case of HMDs, rotational errors can be serious. Displacement errors for the CAVE and monitor suggest that for small dis-

tances between the viewer and the display screen, there is little difference in effect between the CAVE and monitor. For large distances, the angular error is less for the CAVE because of the typically larger distance between viewer and display. For small distances (for example, 20 cm) the monitor has the best performance, and the CAVE and HMD and also BOOM display have slightly worse performance. For large distances (for example, 500 cm), the HMD and BOOM have very good performance, while the CAVE has 2.5 times the error of the HMD and BOOM, and the monitor is worst at 9 times the error.

The CAVE does have some shortcomings. One is cost: the CAVE is large and expensive, although in some applications the cost per person might be more reasonable because it can be shared by a group of users.

2.3.13 University of New Brunswick, Canada

Researchers at the University of New Brunswick are studying the relationship between stereoscopic vision and visual environments responsive to head movements. As previously mentioned, they have experimentally tested the effect of stereo viewing and head coupling on viewers (Ware, Arthur, & Booth, 1993). The display system developed for this work uses an ordinary CRT display and StereoGraphics CrystalEyes shutter glasses. Head tracking was provided by a mechanically linked head tracker, which allowed coupling of the display image to user head movements. There are several advantages to this approach claimed over HMDs. First, as a result of both increased resolution of the monitor display over typical HMD resolution and a decision to reduce the field of view to 30° laterally, the resolution is 2 minutes of arc per pixel rather than the 12 minutes of arc per pixel typical for a HMD. Second, the monitor approach allows presentation of the 3-D images at a depth of field that the viewer can focus on, with images behind or in front of this depth out of focus, which is normal. This is in contrast to HMDs, in which the optics are typically arranged so as to force the viewer to focus at infinity. Third, an error introduced in HMDs when the eyes move is greatly reduced with the head-coupled shutter glass system because the eyes are further from the display. Finally, the monitor is part of the ordinary workspace, allowing (if desired) the user to see the normal desks, tables, and chairs rather than a completely synthetic environment.

The researchers performed two experiments. The first experiment looked at the subjective impression of 3-D images. Here two simple scenes were used that had strong depth cues even when seen without stereopsis. Subjects viewed the scenes in a number of conditions, including with and without stereo vision and with and without head tracking. Subjects were then asked which condition showed the strongest 3-D effect. Perhaps surprisingly, a non-stereo image with head coupling responsive to movement was typically perceived as having the strongest 3-D effect. A second experiment measured performance on a task in which visual depth was necessary and which used the same conditions as the first experiment. Two complex trees were presented in 3-D, with overlapping branches and one leaf tagged. The task required the subject to determine which root the leaf could be traced to. Results indicated that subjects made a much large number of errors for the stereo

only condition (14.7%), compared with the head-coupled monocular condition (3.7%). The condition with both stereo and head-coupled display resulted in the lowest number of errors (1.3%).

In more recent work, the researchers have developed an algorithm that dynamically adjusts the apparent distance between the eyes for a computer-generated, stereoscopic image displayed on shutter glasses (Ware, Gobrecht, & Paton, 1995). The primary motivation for this work was the fact that different scenes appear to be best displayed by use of an eye separation that is different from the "correct" or expected one, given the apparent distance between the viewer and the object viewed. This assertion was tested experimentally by having viewers set their preferred separation for different moving scenes.

2.3.14 University of Washington

An intriguing alternative to the conventional CRT or LCD display is the use of a laser to "directly write" onto the retina. The motivation for attempting this is to increase the resolution and field of view of displays, and to eliminate display screens and heavy imaging optics and thus create a low profile display. The Human Interface Technology Laboratory (HITL) at the University of Washington is developing a prototype retinal display as a long-term project funded by Micro Vision, Inc., of Seattle, Washington, which hopes to manufacture and distribute the device as a product. The project has a series of long-range project goals that are very ambitious. The goals are to create a device that (1) is small and lightweight enough to mount on eyeglasses; (2) has resolution high enough to approach that of human vision; (3) has a large field of view, greater than 100° per eye; (4) has color resolution superior to standard displays; (5) is capable of displaying either in a dedicated or see-through mode; (6) is bright enough for outdoor use; (7) has very low power consumption; and (8) provides a true stereoscopic display.

To meet the acuity of the eye (about 1 minute, that is 0.016° of arc) while also surrounding the user to the maximum extent (based on a single-eye field of view of 135° horizontally and 150° vertically) would require a display with a resolution of about 8000 by 8000 pixels, far beyond the capabilities of today's displays (Holmgren & Robinett, 1993).

In principle, a direct retinal write, or scanned laser, display uses a laser source along with accompanying mirrors or other deflectors to direct a laser beam through the pupil of the eye and onto the retina with a scanning pattern similar to the scanning used in conventional television. Scanning a laser onto a retina has been previously used in scanning laser ophthalmoscopes. The feasibility of building a practical retinal scanner has yet to be demonstrated, and early prototype efforts by the HITL were criticized because of their use of acousto-optic scanners, which would be very cumbersome in a color version because of the need for six different sets of deflectors and lenses (one for the x axis and one for the y axis for each primary color). The primary problem posed is the difficulty of building a scanner that can deflect the beam horizontally.

The HITL/Micro Vision researchers now believe that a practical retinal scanner can be developed using a different approach than that used in their early prototypes. They have developed (and applied for a patent on) a mechanical resonant scanner that has only a single moving part and has a large scan angle and a high scanning frequency. The device is also quite small (0.9 x 1.3 x 2.8 cm), has a uniform and repeatable scan, reflects all colors at the same angle, and is made from common materials at a volume manufacturing cost estimated to be under \$3. A bench-mounted prototype, using the mechanical resonant scanner to scan the beam horizontally, has been developed that provides VGA resolution images (640 points horizontal, 525 lines vertical) in either monochrome or full (RGB) color. The monochrome system uses a red laser diode, while the color system additionally uses green (helium neon) and blue (argon) gas lasers.

One problem with the present scanner is that the beam moves faster at the center of the scan than at the edges, which results in pixels that are wider and brighter in the center than near the edges. This can be corrected by varying the pixel display time and intensity as the beam scans across the image. A second scanner problem is a change in resonant frequency with temperature, which researchers expect to cure with a feedback mechanism that compensates appropriately for temperature variation. The primary disadvantages of the color display are its size and cost, primarily because of the blue and green gas laser sources, since diodes for these colors are not currently available. Work is being done on frequency doubling methods and the use of non-laser sources, including the development of blue and green LEDs.

Future work will concentrate on further development of the mechanical resonant scanner, resolving a problem of the exit pupil size (which is currently quite small), developing methods of generating color with hardware that can fit in a small package, increasing the resolution of the display, and testing for safety. The researchers expect that safety will not be an issue due to the low power of the light sources, but are planning to begin an extensive series of safety tests by an ophthalmologist.

Micro Vision, Inc. intends to enter the market with monochrome displays for applications that need a compact display with minimal power draw and a bright, medium-to-high resolution image. They expect to begin with a small hand-held prototype system that can be used to develop and test applications.

2.3.15 Xenotech, Australia

Researchers at Xenotech in Australia also are working on techniques for 3-D autostereoscopic display. The goal of the effort is to develop systems that have high image quality and that do not require viewing aids such as shutter glasses. They have built a number of prototype systems, with the design of most of them still confidential. However, they demonstrated an advanced prototype in October, 1995 at the Korean Electronics Show in Seoul, Korea.




Figure 26. Xenotech Autostereoscopic Display Overview

The prototype system uses a 30 inch diagonal television monitor and a single viewer sits at a distance of about 1 m in front of the monitor. Using a pair of projectors and a series of mirrors, and a screen with a specialized material, the device presents a different image to each eye. The basic technique is shown in Figure 26, which is from a patent application filed by Xenotech (Richards, 1994). A pair of video projectors are contained in a housing, one for each eye, with the projector for the right eye (14) shown and the projector for the left eye not visible but behind the other projector. The image for the right eye (18) is projected onto two successive mirrors (60 and 62), and then onto a partially silvered mirror (not labeled but that swings on a hinge 68 between the arrows 66 and 70). The image is then reflected off of the partially silvered mirror to a "retroreflective" surface (64) that acts as a mirror, reflecting the image back toward the viewer (72), but passing first through the partially silvered mirror. This image is focused by optics consisting of the partially silvered mirror and the retroreflecting mirror specifically to the right eye of the viewer. A second image is projected in the same way by the second projector, through the same series of mirrors and optics. The retroreflective mirror consists of a special surface with a zigzag pattern at the pixel level, which has the property that an incident beam of light is reflected back at an angle 180° from the incident angle. For example, a beam of light that comes in at a 90° angle to the surface of the retroreflective mirror will be reflected back at a 90° angle, and is thus seen only by one eye and not the other.

A pair of television cameras at the left and right sides of the monitor track the head position and pupil locations of the single viewer. When a viewer moves, the image projected to each eye moves to compensate, either by adjusting the angle of the partially silvered mirror (for movement in the vertical direction) or by lateral movement of a carriage upon which both projectors are mounted (for movements in the horizontal direction).

Xenotech sees the primary advantage of their systems as their use of field sequential video at 50 (or 60 in NTSC) frames per second for each eye to eliminate the flicker that is a problem in competing systems that present video at 25 (or 30) frames per second. Presentation of images at this frame rate is achieved by holding each frame in a buffer memory and presenting it twice. The prototype system operates at standard NTSC or PAL television resolution, but the approach is applicable to much higher resolution. Other claimed advantages are very large screen sizes (the researchers have built a prototype as large as 50 inches diagonal), and very high image brightness. Xenotech has two specific markets in mind for development, these being military applications and the video games market. Current work includes the development of a prototype autostereoscopic display that will allow multiple viewers of the monitor.

2.4 Summary and Expectations

VEs require stereoscopic displays to allow users to perceive 3-D images. The technique used to create stereoscopy is fundamental to its advantages and limitations. The following discusses the major differences between the techniques and summarizes their relative advantages and disadvantages as of the present and as projected in the next five years. This is done for the four basic categories of displays: HMDs, shutter glasses, passive glasses, and autostereoscopic displays.

HMDs are presently the display technology of choice for most VE applications that require a full sense of immersion, primarily because they can allow a relatively wide field of view. At the present time, however, HMDs have severe limitations, primarily resulting from their weight, cumbersome design, low resolution, and limited field of view. There is a trade-off between resolution and field of view: some displays provide a wide field of view but low resolution, while others a narrow field of view with higher resolution. Poor resolution is not always the fault of the display device. In some applications the quality of rendering of complex graphical images is the major factor limiting resolution, though improved graphics hardware can be expected to change this situation in the future. Nonetheless, no HMD today provides a resolution anywhere close to that perceivable by the eye.

At the present time, there are two distinct classes of HMDs: relatively large, expensive, and heavy devices that have a wide field of view and are designed for military and other industrial applications, and the typically much less expensive devices that are small, lightweight, provide a narrow (30° or less in the horizontal plane) field of view, and are designed for the consumer market. Some of the displays intended for the consumer market, which have become available only in the last year or so, are quite inexpensive (\$500 to \$700), appear to be well received by users, and may challenge shutter glasses for the market in inexpensive 3-D displays. Meanwhile, researchers and manufacturers are working to meet users' demand for high resolution, wide field of view displays at low cost. Advances in component technologies are also likely to reduce some HMD problems. In particular, as LCD technology continues to develop, increases in LCD resolution will allow HMDs to achieve a wider field of view, perhaps doubling in resolution and field of view in the next five years. The weight of such HMDs should also decrease somewhat, although the weight of necessary optics is likely to remain a major limiting factor and may be a real barrier.

Will HMDs continue to be used as heavily as they are today? This is difficult to assess. Investigators have shown that a significant immersive effect can be achieved with shutter glasses and projection screens, and even with shutter glasses and a monitor (Ware, Arthur, and Booth, 1993). HMDs may play a role in the too frequent occurrence of simulator sickness, although this relationship is not well understood as yet. Moreover, initial studies have shown that many people have difficulty accommodating their eyes to the infinity position often assumed by HMDs to allow full flexibility in image projection. Such difficulties may contribute to fatigue and encourage the use of other alternative displays where

this is less of a problem. As other technologies improve, the demand for HMDs may decline.

Retinal displays may solve many of the typical HMD problems of weight, cumbersome packaging, brightness, resolution, restricted field of view, cost, and accommodation. While the approach is very challenging, significant progress is being made, and a new method currently being developed may lead to commercial products that display monochrome images within a few years. Color displays are a more difficult problem because of the difficulty of generating blue and green light sources and probably cannot be expected as commercial products within the next five years.

Shutter glasses are a widely used technique for stereoscopic display, having been available for a number of years. This is due to their relative inexpensiveness, lightness, and less cumbersome design compared to HMDs. In the next five years, though, it can be expected that LCD displays used in monitors and projection displays will be widely available as products that switch sufficiently fast to be used with shutter glasses. At the present time, there are only a few manufacturers of shutter glasses, with products being very similar. Shutter glasses do suffer from problems such as crosstalk, but this problem is likely to be reduced in the future as expected advances in LCD technology lead to decreased switching time. Fast LCD displays are being developed for other purposes, including autostereoscopic displays, and a substantial reduction in crosstalk can be expected in the next five years. In addition, techniques for minimizing crosstalk by controlling brightness and timing of image display can be applied. At the present time, shutter glasses are almost always used with CRTs or projection displays, because of the slowness of LCD displays.

Passive glasses have been used very little in VEs. The new technique of using microelectronic fabrication techniques to create polarizing filters at the pixel level does make this technique practical for VEs and may lead to the increased use of passive glasses, particularly for CAVE systems in which multiple images are projected on walls. One advantage of this new passive glasses approach over shutter glasses is reduced flicker, which can be a problem even if a CRT is modified to operate at 120 Hz (resulting in a 60 Hz presentation rate to each eye). Another advantage in a CAVE system is that a single LCD projector can be used rather than multiple CRT projectors (one for each primary color), thus reducing the cost and amount of maintenance required to keep the projectors in proper alignment. While CRT and shutter glass systems presently have higher resolution than the polarization at the pixel level approach, it is expected that improved fabrication methods will allow CRT-like resolution with the latter method within the next few years.

A considerable number of different autostereoscopic display systems have been developed to the prototype stage. In particular, a significant amount of research effort is being invested in autostereoscopic technology by large Japanese display, electronics, and telecommunications companies, including Sharp, Canon, ATR, Sanyo, and Teuromo. In addition, many small start-up companies have been recently formed to develop autostereoscopic systems. As yet only one commercial product is being sold (the Virtual Window dis-

play by Dimension Technologies, Inc.) which can be expected to be used in VE applications. This is expected to change with several additional products coming to market within the next few years. Increased resolution of flat panel displays and a trend toward the display of larger numbers of perspective views simultaneously will substantially reduce the current major limitation of these systems, which is that users are limited in their lateral movement.

The relatively large amount of research and development in autostereoscopic display technology may seem surprising since there is no proven market. Different developers of autostereoscopic displays have different motivations, with some more concerned with games and military and industrial applications, and others more concerned with 3-D television. However, the high costs and the difficulty on agreeing on standardization suggest that certain applications, possibly including games, military, and industrial applications will be the initial markets, with 3-D television coming later. The use of these displays in VEs is dependent upon the (lack of) acceptance of HMDs and shutter glasses by users, as well as the type of application involved.

3. TRACKING INTERFACES

Tracking, also called Position and Orientation Tracking or Position Tracking and Mapping, is used in VEs where the orientation and the position of a real physical object is required. Specifying a point in 3-D requires the transition position, that is, the Cartesian coordinates x , y , and z . However, many VE applications manipulate entire objects and this requires the orientation to be specified by three angles known as pitch (elevation), roll, and yaw (azimuth). Thus, six degrees of freedom (DOF) are the minimum required to fully describe the position of an object in 3-D.

Trackers are used to measure the motion of the user's head or hands, and sometimes eyes. This information is then used to correlate visual and auditory inputs to the user's position. In this way, trackers are one part of a visually coupled system that Kocian and Task (1995) define as a special subsystem integrating the natural visual and motor skills of an operator into the system he is controlling. For example, in the case of magnetic sensors, a receiver is placed on the user's head so that when the head moves, so does the position of the receiver. The receiver senses signals from the transmitter which generates a low frequency magnetic field. The user's head motion is sampled by an electronic unit that uses an algorithm to determine the position and orientation of the receiver in relation to the transmitter. In addition to magnetic head trackers there are mechanical, optical, acoustic (ultrasonic), and inertial head trackers. These types of trackers also can be mounted on glove or body suit devices to provide tracking of a user's hand or some other body part, see Section 5.1. Some include special facilities to augment the tracking function with 3-D mouse-like operations. Eye trackers work somewhat differently; they do not measure head position or orientation but the direction at which the users' eyes are pointed out of the head. This information is used to determine the direction of the user's gaze. Eye trackers use electroocular, electromagnetic, or optical technologies.

Trackers also are used in augmented reality applications. In these systems, the user sees the real world around him with computer graphics superimposed or composited with the real world. One of the big obstacles to widespread use of artificial reality is the registration problem in correctly aligning real and virtual objects. Because of lags in the time interval between measuring the head location and superimposing the corresponding graphic images on the real world, virtual objects may appear to swim around real objects. Since the human eye is very good at detecting even very small misregistrations, errors that can be tolerated in immersive VE are not acceptable in augmented reality, though registration is a tracking issue for all types of applications when multiple moving objects are involved.

This section starts by discussing trackers under the heading of head tracking although, as indicated above, many of these trackers also can be used for hand and body tracking. It then moves on to look at eye tracking. The final part of the section presents some projections for how tracking capabilities, as a whole, may evolve over the next few years.

3.1 Head Tracking

Several researchers have investigated the value of head tracking for promoting a sense of immersion in VEs. One study compared head tracking and hand tracking for an experimental task that required subjects to visual scan a room and locate targets consisting of two-digit numbers; in both cases, the method of tracking was used to control the visual scene displayed (Pausch, Shackelford, and Proffitt, 1993). The results of this experiment showed that subject using head tracking were nearly twice as fast in located targets as those using hand tracking (a mean of 1.5 seconds per target for head tracking versus 2.6 seconds per target for hand tracking). Hendrix (1995) reports that the use of head tracking, together with a stereoscopic visual display, significantly increased the reported sense of presence in a VE and subjects' subjective assessment of their performance in performing spatial judgments, although actual performance measures showed no increase in judgement accuracy.

As previously stated, there are several tracking technologies in use, although the most common are magnetic, acoustic, and optical technologies. Table 3 provides an overview of the different technologies and their advantages and disadvantages.

Head trackers can be described with respect to a small set of key characteristics that serve as performance measures for their evaluation and comparison. Meyer et al (1992) define these characteristics as resolution, accuracy, and system responsiveness (additional characteristics of robustness, registration, and sociability are not considered here).

- Resolution. Measures the exactness with which a system can locate a reported position. It is measured in terms of inch per inch of transmitter and receiver separation for position, and degrees for orientation.
- Accuracy. The range within which a reported position is correct. This is a function of the error involved in making measurements and often it is expressed in statistical error terminology as degrees root mean square (RMS) for orientation and inches RMS for position.
- System responsiveness. Comprises:
 - Sample rate. The rate at which sensors are checked for data, usually expressed as frequency.
 - Data rate. The number of computed positions per second, usually expressed as frequency.
 - Update rate. The rate at which the system reports new position coordinates to the host computer, also usually given as frequency.

Table 3. Tracking Technologies

Technology	Description	Strengths	Weaknesses
Mechanical	Measure change in position by physically connecting the remote object to a point of reference with jointed linkages	Accurate Low lag No line of sight (LOS) or magnetic interference problems Good for tracking small volumes accurately	Intrusive, due to tethering Subject to mechanical part wear-out
Magnetic	Use sets of coils (in a transmitter) that are pulsed to produce magnetic fields. Magnetic sensors (in a receiver) determine the strength and angles of the fields. Pulsed magnetic field may be AC or DC.	Inexpensive Accurate No LOS problems Good noise immunity Map whole body motion Large ranges - size of small room	Ferromagnetic and/or metal conductive surfaces cause field distortion Electromagnetic interference from radios Accuracy diminishes with distance High latencies due to filtering
Sourceless, Non-inertial	Use passive magnetic sensors, referenced to the earth's magnetic field, to provide measurement of roll, pitch, and yaw, and as a derivative, angular acceleration and velocity.	Inexpensive Transmitter not necessary Portable	Only 3 DOF Difficult to mark movement between magnetic hemispheres
Optical	Use a variety of detectors, from ordinary video cameras to LEDs, to detect either ambient light or light emitted under control of the position tracker. Infrared light is often used to prevent interference with other activities.	High availability Can work over a large area Fast No magnetic interference problems High accuracy	LOS necessary Limited by intensity and coherence of light sources Weight Expensive
Acoustic (Ultrasonic)	Use three microphones and three emitters to compute the distance between a source and receiver via triangulation. Use ultrasonic frequencies (above 20 kHz) so that the emitters will not be heard.	Inexpensive No magnetic interference problems Lightweight	Ultrasonic noise interference Low accuracy since speed of sound in air varies with environmental conditions Echoes cause reception of "ghost" pulses LOS necessary
Inertial	Use accelerometers and gyroscopes. Orientation of the object is computed by jointly integrating the outputs of the rate gyros whose outputs are proportional to angular velocity about each axis. Changes in position can be computed by double integrating the outputs of the accelerometers using their known orientations.	Unlimited range Fast No LOS problems No magnetic interference problems Senses orientation directly Small size Low cost	Only 3 DOF Drift Not accurate for slow position changes

- Latency, also known as lag. The delay between the movement of the remotely sensed object and the report of the new position. This is measured in milliseconds (ms).

Another pertinent characteristic is repeatability. The accuracy of the measured 3-D position and orientation between bodies is based on the composite effects of the measurement separability (that is, variance) and the measurement of offset (that is, measurement errors that cannot be removed by collecting more data). Thus, repeatability refers to the resulting distribution spread of several repeated measurements of, for example, a single stationary point. It provides a gauge of measurement precision and can be expressed in inches, degrees, or microns. Drift is problem specific to inertial trackers. These trackers combine measurements from accelerometers and gyroscopes to calculate 6 DOFs relative to the starting position. Since the measurement is relative, rather than absolute, drift can cause errors in reading that require an absolute position for re-calibration.

These characteristics provide some guidance for tracker performance. One of the most important is latency. Durlach (1994) states that delays greater than 60 msec between head motion and visual feedback impair adaptation and the illusion of presence. Latencies of greater than 10 msec may contribute to simulator sickness. Bryson (1993) considers systems with latency longer than 0.5 second not to be real-time interactive. On the other hand, in the case of non-immersive VEs systems where the VE is viewed through a CRT, Ware and Balakrishnan (1994) found latency in the head-tracking system to be relatively unimportant in predicting performance, whereas latency in the hand-tracking system was critical. Latency between systems are difficult to compare because they are not always calculated the same. Bryson (1993) identifies several sources of latency: delays in the tracker signal, delays in communication between the tracker and the computer system, delays due to computations required to process the tracker data, and delays due to graphical rendering. However, several manufacturers contacted for this report suggested that 1/frequency as the preferred measure.

With respect to responsiveness, Durlach (1994) contends that head movements can be as fast as 1,000°/sec in yaw, although more usual peak velocities are 600°/sec for yaw and 300°/sec for pitch and roll. The frequency content of volitional head motion falls off approximately as $1/f^2$, with most of the energy contained below 8 Hz and nothing detectable above 15 Hz. Tracker to host reporting rates must, therefore, be at least 30 Hz.

Where the information was available, data on resolution, accuracy, and system responsiveness are given for all the commercial tracking products described below. An additional important characteristic that is included is working volume or range, which may be bound by intrinsic limitations such as mechanical linkage or signal strength. This is the volume in which a position tracker accurately reports position. It is variously expressed in feet or meters, inches or feet in diameter, or as some portion of a geometric shape such as a sphere.

Table 4. Characteristics of Commercially Available Tracking Devices

Technology	Product	Vendor	DOF	Frequency	Latency	Resolution	Working Volume	Price
Mechanical	BOOM3C	FakeSpace, Inc.	6	>70 Hz	200 msec	0.1°	6 ft diameter, 2.5 ft high	\$95,000 ^a
	PUSH	FakeSpace, Inc.	6	>70 Hz	200 msec	0.1°	2 ft diameter	\$45,000 ^a
	ADL-1	Shooting Star Technology	6	240 Hz	0.35 - 1.8 msec	0.25 in, 0.15 - 0.3°	35 in diameter, 18 in high	\$1,299
	WrightTrac	Vidtronic, Inc.	6	300 Hz	3.3 msec	0.1° per axis	1/4 sphere, 40 in diameter	\$795
Magnetic	Fastrak	Polhemus	6	120 Hz/number of receivers	4 msec ^b , 8.5 msec ^{bc}	0.0002 in/in, range 0.025°	10 - 30 ft	\$6050
	Isotrak II	Polhemus	6	60 Hz/number of receivers	20 msec ^b , 40 msec, 23-45 msec ^b	0.0015 in/in, range 0.1°	5 ft	\$2,875
	Insidettrak	Polhemus	6	60 Hz/number of receivers	12 msec ^b	0.0003 in/in, range 0.03°	5 ft	\$999
	Ultratrak	Polhemus	6	30 - 60 Hz	12 msec ^b	0.05 in at 5 ft range, 0.25 ft at 15 ft range 0.1° RMS	2 - 15 ft	\$23,250 - \$32,250
	Ultratrak 120	Polhemus	6	60 - 120 Hz	12 msec ^b	0.15 in at 5 ft range, 0.25 ft at 15 ft range 0.1° RMS	2 - 15 ft	\$39,500 - \$71,500
	Flock of Birds	Ascension Technology Corp.	6	Up to 144 Hz	7.5 msec ^b , 39-47 msec	0.1° RMS at 12 in	3 - 8 ft	\$2,695 basic system, \$8,090 ERT system
	PC/BIRD	Ascension Technology Corp.	6	Up to 144 Hz	10 msec	0.08 in translation, 0.15° rotation	4 ft	\$2,475
Sourceless, non-inertial	SpacePad	Ascension Technology Corp.	6	120 Hz/number of receivers	<8 msec ^c	Not available	16 x 16 ft	\$985
	CyberTrack 3.2	General Reality Company	3	30 Hz	<50 msec	0.15° heading, 0.12° tilt	360° horiz, ±55° tilt	\$850
	Wayfinder - VR	Precision Navigation, Inc.	3	30 Hz	33.3 msec ^d	±0.1° heading, ±0.1° tilt	360° horiz, ±20°, ±50° tilt	\$599 - \$699

Table 4. Characteristics of Commercially Available Tracking Devices

Technology	Product	Vendor	DOF	Frequency	Latency	Resolution	Working Volume	Price
Sourceless, non-inertial	CyberMaxx	Victor Maxx Technologies	3	75 Hz	29.6 msec	0.1° vertical, 0.1° horiz	360° horiz, ±45° tilt	\$799
	Mouse-Sense 3D	RPI Advanced Technology Group	3	8 Hz	125 msec ^c	±1° heading, ±1° tilt	360° horiz, ±25 - 55° tilt	\$750
Optical	SELSHOT II	Selcom AB	6	10 KHz	Not given	0.025% of millirad	Up to 200 m	\$29,980
	OPTOTRAK 3020	Northern Digital, Inc.	6	600 Hz	Not given	0.01 mm at 2.25 m	Not given	\$57,400
	MacReflex	Qualisys, Inc.	6	50-200 Hz	Not given	Not given	0.5 - 30 m (indoors), 0.5 - 9 m (outdoors)	\$38,500 - \$48,500
	DynaSight	Origin Instruments Corporation	3	64 Hz max	16-31 msec	0.1 mm cross range 0.4 mm down range	0.1 - 1.5 m for 7 mm target, up to 1 - 6 m for 7 mm target	\$2,195
	RK-447 Multiple Target Tracking System	ISCAN, Inc.	6	60 Hz	16 msec	Not given	Not given	\$36,800
Acoustic (Ultrasonic)	Head/Hand XYZ Tracker	Fifth Dimension Technologies	3	20 Hz	Not given	Not given	Up to 3 m	\$345
	GP12-3D (Free-point 3D)	Science Accessories Corp.	3	150 Hz/number of emitters	Not given	0.002 in	3.25 x 3.25 x 3.25 ft up to 16 x 8 x 8 ft	\$4,995 - \$6,995
	Logitech 3D Mouse and Head Tracker	Logitech	6	50 Hz	72 msec ^c	1/250 in (linear), 1/10° angular	5 ft long, 100° cone	Not given
Inertial	MotionPak	Systron-Donner	6	60 Hz	Not given	0.004°/sec.	Not given	\$10,000
	GyroPoint	Gyration, Inc.	3	Not given	Not given	0.2°	75 ft	\$299

a. Includes visual system.

b. Unfiltered

c. (Adelstein, Johnston, and Ellis, 1995)

d. Calculated as 1/Frequency and converted to milliseconds.

e. (Ware and Balakrishnan 1994).

3.1.1 Commercially Available Trackers

This section presents over twenty five different tracking devices. Some tracking systems are integral parts of a display system and these are described in the Section 2.2 of this report. Specifically, the Fakespace, Inc. BOOM 3C and PUSH are described in Section 2.2.14 and Section 2.2.17, respectively; and the VictorMaxx Technologies CyberMaxx is described in Section 2.2.4. Additional products, Sensor Applications Inc.'s CG93 and SPAR non-inertial sourceless sensors and MVR, Inc.'s Optical Head Tracker could not be included for lack of adequate information. Table 4 presents summary information for all the described devices, except BioVision, by Optimum Human Performance Center, Mandala Virtual Reality Systems from the Vivid Group, and REALWare VR System by CCG Meta-Media, Inc. These were included in the text, but excluded from the table for lack of specific tracking information.

In addition to the following products, Crossbow Technology has recently announced the availability of a new tracking technology. Designed for use in 3-D games, this technology employs small, low-cost, silicon micromachined accelerometers and silicon-based magnetic sensors (a patent is pending on these microsensors). It supports roll, pitch, and yaw tracking with a resolution of less than 0.1° , and a speed of less than 10 msec. Crossbow Technology has demonstrations available for a 3-axis orientation tracker, called the TRK300, and a 2-axis analog joystick system. The company develops custom designs to meet specific user needs, and currently is looking for OEMs to license the base tracking technology.

3.1.1.1 ADL-1

ADL-1 by Shooting Star Technology is a 6 DOF mechanical head tracker. The user wears a headband attached to the lightweight arm while seated before a video display for non-immersive VR, or can attach the tracker to a HMD for conventional immersive VR. Sensors mounted on the arm measure the angles of the joints of the arm. A microprocessor uses the angles to compute the head's geometry and sends the data to a host computer via a serial connection. Specification details for the ADL-1 are given in Figure 27. ADL-1 costs \$1,299.

Specification	
Update Rate	Max. 240 Hz
Sampling Rate	240 Hz
Latency	0.35-1.8 msec
Accuracy	0.2 in
Linear Resolution	~ 0.025 in
Angular Resolution	$0.15 - 0.3^\circ$
Work Volume	Half cylinder, ~ 36 in diameter, 18 in height
Repeatability	0.1 in

Figure 27. ADL-1

3.1.1.2 Vidtronics Wrihtrac

Specification	
Update Rate	300 Hz
Resolution	0.1°
Work Volume	1/4 sphere, ~ 40 in diameter.
Repeatability	0.1 in

Figure 28. Wrihtrac

Wrihtrac, by Vidtronics, Inc. is a 6 DOF mechanical tracker designed for use with PCs and desktop VR systems. It consists of an aluminum arm and boom. It has microprocessor-based control and a potentiometer-based position/orientation system. The Wrihtrac uses serial RS-232 data transmission and comes with a DOS interface. Some specification details are given in Figure 28. Wrihtrac costs \$795.

3.1.1.3 Fastrak

The Polhemus Fastrak was developed based on a redesign of the vendor's Isotrak product. Isotrak was one of the first trackers developed for use in VEs. It used older analog technology that produced large latencies, sensor range was small, and signal noise large. These problems were solved in the redesign by using a digital signal processing architecture.

Each Fastrak can accept data from up to 4 receivers and up to 8 systems can be multiplexed to allow up to 32 receivers. Further details are given in Figure 29. Fastrak costs \$6,050.

Specification	
Update Rate	120 Hz + no. receivers
Latency	4 msec
Linear Accuracy (RMS)	0.03 in
Angular Accuracy (RMS)	0.15°
Resolution	0.0002 in/in, 0.025° orientation
Range	Up to 10 ft, 30 ft with Long Range transmitter

Figure 29. Fastrak

3.1.1.4 Isotrak II

The Polhemus Isotrak II is a lower cost tracking system with slightly reduced performance from Fastrak. It consists of an electronics unit, a single transmitter, and 1 or 2 receivers. One of these receivers can be the Polhemus 3Ball, a 3-D positioning/orientation input device that incorporates hand and body motion into a mouse. This optional receiver costs

\$800 in addition to \$2,875 for the basic Isotrak II. Specification details for the Isotrak II are given in Figure 30. In an independent study of lag and frame rates of VE systems, Ware and Balakrishnan (1994) found the device lag for the Isotrak II to be 45ms.

Specification	
Update Rate	60 Hz + no. receivers
Latency	20 msec (unfiltered)
Accuracy (RMS)	0.1 in for x, y, and z, 0.75° for receiver orientation
Resolution	0.0015 in/in, 0.1° orientation
Range	Up to 15 ft

Figure 30. Isotrak II

3.1.1.5 Insidettrak

Specification	
Update Rate	60 Hz + no. receivers
Latency	12 msec (unfiltered)
Accuracy (RMS)	0.5 in x, y, and z, 2° for receiver orientation; < 30 ft range
Resolution	0.0003 in/in, 0.03° orientation
Range	Up to 5 ft

Figure 31. Insidettrak

Insidettrak, by Polhemus, is an even smaller version of the Fastrak sensor, compressed to fit on a PC-card that plugs into 386 or 486 PCs. Specification data for Insidettrak are given in Figure 31. Testing by Burdea and Coiffet (1994) found that Insidettrak sensing data is much noisier than Fastrak. In June

1995, Polhemus announced a drop of price for Insidettrak from \$2,250 to \$999.

3.1.1.6 Ultratrak

The most expensive of the Polhemus offerings is the newly introduced Ultratrak, an integrated motion capture system designed to meet the needs of applications requiring full-body motion tracking.

Ultratrak consists of a 486-based Motion Capture Server unit which contains 4 to 8 motion capture boards (each board can support 2 receivers), a VGA controller, external synchronization board, and communications card. The base system comes with 8 receivers and up to an additional 8 receivers may be added to the system. Moreover, multiple systems can be networked. Ultratrak comes in a 60 Hz version and a 120 Hz version (Ultratrak 120). Both come with the Long Ranger transmitter (optional equipment for Fastrak and Insidettrak) that allows tracking and capturing a subject in an area in excess of 700 square feet. Further details for Ultratrak are given in Figure 32. Ultratrak costs \$23,250 for an 8 receiver system and up to \$32,250 with all 16 receivers. An Ultratrak 120 costs between \$39,500 and \$71,500, depending upon the number of receivers.

Specification	
Update Rate	60 Hz up to 8 receivers, 30 Hz up to 16 receivers
Latency	20 msec (unfiltered)
Linear Accuracy (RMS)	0.1 in at 5 ft, 3 in at 15 ft
Angular Accuracy (RMS)	0.75°
Resolution (RMS)	0.005 in at 5 ft, 0.25 in at 15 ft; 0.1° orientation
Range	Up to 15 ft

Figure 32. Ultratrak

3.1.1.7 Flock of Birds

Flock of Birds is a 6 DOF tracking system by Ascension Technology Corporation. It is intended for tracking human motions in character animation, biomedics, and VE applications. In particular, Flock trackers are used for head tracking in flight simulators/trainers; head, hand, and body tracking in VE

games; and full body tracking for character animation, performance animation, virtual walkthroughs, and sports analysis. Flock of Birds has full 360° coverage without blocking or echoing problems and a fast measurement rate—up to 144 position and orientation measurements per second. It can simultaneously track up to 60 separate independent points out to 8 feet with the Extended Range Transmitter option. Each standard-range transmitter allows operation in about a 3-foot radius. Ascension claims it has the lowest lag of all trackers when tracking multiple points. Specification details can be found in Figure 33. The Flock of Birds emitter radiates a sequence of DC pulses, in effect switching the emitted field off and on. This design is intended to reduce the effect of distorting eddy currents

Specification	
Update Rate	Up to 144 Hz
Angular Accuracy (RMS)	0.5°
Translation Accuracy (RMS)	0.1 in
Angular Resolution (RMS)	0.1° at 12 ft
Translation Resolution	0.03 in
Translation Range	±3 ft (8 ft optional) any direction
Angular Range	±180° yaw, roll; ±90° pitch

Figure 33. Flock of Birds

induced by changing magnetic fields in metallic objects. While it minimizes the effect of conductive metals, the Flock of Birds remains sensitive to ferromagnetic metals.

A Flock of Birds system with one receiver costs \$2,695. With the Extended Range Transmitter, the system costs \$8,090. Additional receivers for either configuration cost \$2,245.

3.1.1.8 PC/BIRD

PC/BIRD is a new offering from Ascension Technology Corporation that uses the same patented pulsed-DC magnetic technology employed in the other Ascension tracking products. Intended for use with PCs, this tracker is configured as an ISA-compatible board, a receiver that can be mounted on any non-metallic object, and either a standard or extended range transmitter. With the standard range transmitter, PC/BIRD operates with a range of 4 feet, the extended range transmitter allows a range of up to 10 feet. Measurements are made at the rate of up to 144 per sec. Additional cards and receivers may be used to track multiple objects simultaneously. Further details are given in Figure 34.

An optional mouse, with three programmable buttons, is available for providing user inputs in 2-D or 3-D. The list price for the basic PC/BIRD is \$2,475. The extended range transmitter cost is \$5,845, and the 3D mouse option, in lieu of a receiver in a shell, is \$345.

3.1.1.9 SpacePad

Another recent product from Ascension Technology Corporation is a low-cost magnetic tracker, SpacePad, intended for use by VE game developers and designers of interactive experiences. SpacePad measures the position and orientation of one or more lightweight receivers attached to a person. SpacePad makes 120 measurements per second in its single receiver mode (up to four receivers can be used). Lag is less than 8ms, as shown in Figure 35. Since the SpacePad is intended for use in an immersive environment, Ascension considers accuracy and resolution to be less important than update rate and lag. Consequently, Ascension has not calculated those parameters. Range is configuration dependent; the larger the antenna loops (up to 16 x 16 feet), the greater the tracking volume. A single-receiver board set costs around \$985.

Specification	
Update Rate	Up to 144 Hz
Latency	10 msec
Translation Accuracy	0.08 in
Angular Accuracy	0.15°
Range	Up to 4 ft, 10 ft with extended range transmitter

Figure 34. PC/BIRD

Specification	
Update Rate	120 Hz for 1 received, 60z for 2 receivers 30 Hz for 4 receiver, s
Latency	8 msec
Translation Range	Configuration dependent
Angular Range	±180° azimuth, roll, ±90° elevation
Range	Up to 16 x 16 ft

Figure 35. SpacePad

3.1.1.10 CyberTrack 3.2

Specification	
Sampling Rate	Max 30 Hz
Latency	<50 msec
Accuracy	$\pm 1.25^\circ$ heading, $\pm 0.25^\circ$ tilt
Resolution	0.15° heading, 0.12° tilt
Range	360° horizontal, $\pm 55^\circ$ tilt
Repeatability	0.25° heading, 0.25° tilt

Figure 36. CyberTrack 3.2

3.2 costs \$850.

3.1.1.11 Wayfinder-VR

Precision Navigation, Inc.'s Wayfinder-VR is another low-cost, sourceless head tracker. It is a passive attitude detection system based upon a proprietary triaxial magnetometer system and a biaxial electrolytic inclinometer. It uses a 3-axis magnetometer to sense the earth's magnetic field and a 2-axis tilt sensor to measure pitch and roll. It combines these data to mathematically compute orientation and output heading, pitch, and roll data via RS-232 serial link. In addition, a mouse emulation mode that maps yaw to left-right motion and pitch to up-down motion is available. Specification details are given in Figure 37. The device costs between \$599 and \$699 depending on the tilt range required.

Specification	
Sampling Rate	Max 30 Hz
Accuracy	$\pm 2^\circ$ heading, $\pm 1^\circ$ tilt
Resolution	0.1° heading and 0.1° tilt
Range	360° horizontal, $\pm 20^\circ$, $\pm 50^\circ$ tilt
Repeatability	1° heading and 0.5° tilt

Figure 37. Wayfinder-VR

3.1.1.12 Mouse-Sense3D

Specification	
Update Rate	8 Hz (higher rates available as custom option)
Accuracy	$\pm 1^\circ$ heading, $\pm 2^\circ$ tilt
Resolution	$\pm 1^\circ$ heading, $\pm^\circ$ tilt
Tilt Range	$\pm 25^\circ$ ($\pm 55^\circ$ option available)

Figure 38. Mouse-Sense3D

are 2.375 x 4.25 inches. Further details are given in Figure 38.

RPI, Advanced Technology Group, has produced the Mouse-Sense3D. This product is described as a low-cost (\$750.00), high-end, multi-use position sensor for head tracking, body tracking, and three-space gesture tracking. It weighs 2.75 ounces and its dimensions

3.1.1.13 Selcom AB, SELSPOT II

SELSPOT II is a commercial tracking system marketed by Selcom AB, a Swedish company. A camera registers light pulses from LEDs attached to the object being tracked. Located between the lens and electronics of the camera is the SELSPOT sensor, a patented photodetector made by SiTek Laboratories and consisting of a flat semi-conductor disc. Each side of the diode has a light-sensitive coating to produce a high resolution, two-axis field. When a light pulse from one of the LED's passes the lens system in the camera and strikes a point within this field, the electronics registers the x and y coordinates in the two-axis field. Two or more cameras are required to analyze movements in three dimensions.

Specification	
Sampling Rate	10 kHz
Resolution	0.025% of milliradians
Number of Cameras	Up to 16
Number of LEDs	Up to 120
Range	Up to 20 m, 200 m with LED 9

The motion analysis system is capable of analyzing two or three dimensional motion in real time. Specification details are given in Figure 39. Prices start at \$29,980.

Figure 39. SELSPOT II

3.1.1.14 OPTOTRAK 3020

The OPTOTRAK 3020 by Northern Digital Inc. is an infra-red (IR)-based, non-contact position and motion measurement system. Small IR LEDs (markers) attached to a subject are tracked by a number of custom designed sensors. The 3-D positions of the markers are determined in real-time or post hoc, up to 256 markers can be tracked. The

Specification	
Max. Data Rate	3500 Hz (raw), 600 Hz (real-time 3-D)
Accuracy (RMS)	0.1 mm for x, y and 0.15 mm for z, at 2.25 distance
x, y Resolution	0.01 mm at 2.25 m distance
Max. Markers	256
FOV	1.28 x 1.34 m at 2.25 m distance, 2.6 x 3.54 m at 6 m distance

Figure 40. OPTOTRAK 3020

position sensor consists of three 1-D charged coupled device (CCD) sensors paired with three lens cells and mounted in a 1.1m long stabilized bar. Within each of the three lens cells, light from the LED is directed onto a CCD and measured. All three measurements together determine the 3-D location of the marker, which is calculated and displayed in real time. Specification details are given in Figure 40.

The standard OPTOTRAK 3020 system includes one position sensor unit, a kit of 24 markers, a system control unit, standard data collection, display, and utility software, together with cables and other hardware. It costs \$57,400. Additional sensors are \$47,500 each and up to 8 position sensors can be used per system.

3.1.1.15 MacReflex Motion Measurement System

The MacReflex Motion Measurement System, by Qualisys, Inc. also is designed to measure the 3-D motion of subjects in real-time. The system is comprised of (1) one or more MacReflex position sensors (a 3-D system uses from two to seven position sensors), (2) software to enable the

Specification	
Sampling Rate	50-200 Hz
Max. Markers	20
Range	0.5-30 m indoors, 0.5-9 m outdoors

user to set up and calibrate the field of view of the position sensors, and process the measured spatial coordinates of the target markers that are attached to the subject being tracked, (3) passive reflective target markers, (4) a calibration frame for 3-D measurements, and (5) a Macintosh computer system. The position sensor has two components: a CCD digital video camera, and a video-processor. The camera views up to 20 markers in real-time. It then sends the video image to the video processor which determines the centroid of each marker and determines its x, y coordinates. A program converts the x, y coordinates to enable cal-

Figure 41. MacReflex

ulation of position, displacement, velocity, acceleration, angles, angular velocity, and angular acceleration. Some specification details are given in Figure 41.

A complete 60 Hz system costs \$38,500. A 120 Hz system costs \$48,500. Additional position sensors are \$13,500 and \$17,500 for 60 Hz and 120 Hz, respectively.

3.1.1.16 DynaSight

The Origin Instruments Corporation tracking product, DynaSight, is an electro-optical sensor with integrated signal processing that performs 3-D measurements of a passive, non-tethered target. A two-color LED on the front of the sensor indicates the tracking status to the user. In a typical application, the sensor is mounted just above the viewable area of a real-time graphics display. The sensor's field of view is a nominal 75° cone, and the sensor is pointed such that this field covers the comfortable range of head/eye positions for the user of the display. The sensor measures and reports on the 3-D movements of a tiny target that is referenced to the user's forehead. The passive target itself can be mounted on eye glasses, stereoscopic goggles, or on the user's forehead. Larger high-performance targets are available that allow measurements at a sensor-to-target range of up to 20 feet.

	Specification
Update Rate	Max. 64 Hz
Latency	16-31 msec (operating mode dep.)
Resolution (RMS)	0.1 mm cross range, 0.4 mm down range
Accuracy (RMS)	2 mm cross range, 8 mm down range
Lock-on Delay	0.3 sec
Range	0.1-1.5 m for 7mm target, 0.3-4 m for 25 mm target, 1-6 m for 75 mm target

Figure 42. DynaSight

The Active Target Adapter enables tracking of up to four active targets tethered to the Adapter. Five DOF are achieved with two targets, while 6 DOF can be achieved by tracking three or four active targets.

DynaSight is the first in a new line of 3-D measurement products. It is planned that future systems will offer 6 DOF for HMDs

using passive sensors and multiple sensors for networked operations in large virtual volumes. Specification details for DynaSight are given in Figure 42. (In this Figure, measurement parameters for resolution and accuracy are quoted for a 7 mm target at 80 cm range under normal fluorescent room lights.) The product cost is \$2,195.

3.1.1.17 BioVision

Optimum Human Performance Centres, Inc. market a product designed to support animation. Called BioVision, this system uses multiple high speed cameras and lightweight retroreflective markers to capture motion at 60-200 frames per second. Motion is digitized, producing 3-D coordinates for each marker for each frame of the motion and software provides 3-D position, rotation, and scaling information for each of the body parts. The digitized data can then be viewed on a Silicon Graphics workstation or a PC running 3D Studio. BioVision provides 6 DOF information for each body part, but not in real-time. Currently, the motion capture system can go up to 25 markers, which is generally enough to cover one person. The next generation system is expected to be able to handle two people through higher resolution cameras and software that can manage up to 100 markers.

Like all optical systems, occlusion is a problem; however, since BioVision is not a real-time system, occluded markers can be edited. Other special features include the ability to edit bad or missing data, the ability to merge files together, and the ability to export data to all the major software animation packages. The system is marketed as a product, but price information is not available. More typically, a client can purchase BioVision services on a daily basis. Motion capture fees are \$2,700 per day for one to three days; \$2,400 per day for four to six days, and \$2,200 per day for seven days and more. Processing fees, which includes tracking, editing, motion conversion, and data quality review cost \$1,600 per man-day for the first twenty days and \$1,300 per man-day thereafter.

3.1.1.18 Mandala Virtual Reality Systems

The Mandala Virtual Reality Systems from the Vivid Group use computer vision with video cameras as an input device to allow for motion tracking. Their software library contains a complete line of sports applications and other arcade adventures. An example is Turbo Kourier, a flying experience that allows a player to guide their Skyboard through the obstacles of a futuristic city while gathering valuable packages. Mandala offers four systems: the Mandala Virtual Reality System, the Mandala Standard Touring Unit, the Mandala Promotional Touring Unit, and the Mandala VR Module. The systems range in cost from \$21,000 to \$29,000.

All systems include cameras, a CPU, hard drive, specialty cards, and VGA monitor plus a choice of 4 Mandala virtual worlds (1 feature and 3 attractions). The Mandala VR module, their newest product consists of a booth (9.8 feet high, 10 feet wide, and 9.08 feet deep) which houses a virtual stage, speakers, lights, monitors, and camera.

3.1.1.19 REALWare

The REALWare VR system by CCG MetaMedia, Inc. also supports unencumbered VEs. The player interacts with the VE by wearing a colored cotton glove. A video camera focuses on a chromakeyed player standing before a blue wall and inserts the player's image into the VE, which can appear on everything from a TV monitor or a projection system, to a video-wall. As the player moves, a computer tracks the colored glove and reacts to its motion.

REALWare runs on two 486-based PCs, one for simulation and video control and the other for tracking. The optical tracking system returns the location of the centroid of the user's gloved hand 30 times per second. Participants are scanned at the beginning of a simulation to determine the colors of their clothing. Then the system selects a glove color that has the least chroma/luminance overlap with the clothing colors. Color calibration is fine-tuned in a 30 second procedure in which the participant touches a series of virtual objects. REALWare applications include Virtual Hoops (a basketball game) and T-probe (Virtual Voyage to XIA), a multilevel game. The integrated package (hardware and software) costs \$35,000.

3.1.1.20 RK-447 Multiple Target Tracking System

The RK-447 Multiple Target Tracking System, by ISCAN, Inc., is a video tracking system which can track up to 64 facial points at 60 Hz with a latency of 16 msec. It is a real time digital image processor employing ISCAN's proprietary Simultaneous Multiple Area Recognition and Tracking (SMART) architecture. The ISCAN SMART processor computes the position and size of up to 256 areas that are within a particular range of intensity levels. Filtering the output of the SMART processor allows the complete system to specify targets of desired size, position, and intensity parameters from a field containing many potential targets.

After positioning the imaging sensor to include the desired field of view, the image gray level corresponding to the target may be selected. The areas of the video image whose intensity is within the gray level threshold setting are presented on the monitor as a bright overlay, letting the operator see precisely the video information being processed. For each thresholded area, size and position data are computed and stored in a data table which may be accessed by an external computer.

The RK-447 Multiple Target Tracking System divides the image signal into a 512 horizontal by 256 vertical picture element matrix. As the targets' position and size data are automatically determined over the monitor image area, the data within the azimuth and elevation coordinate table correspond to the horizontal and vertical coordinates within the video matrix. These coordinate data are updated every 16 msec and are available for input to a computer. Parametric information may be input to the RK-447 to automatically limit the data set to targets within a particular size or position range. The system costs \$18,500.

3.1.1.21 Head/Hand XYZ Tracker

The Fifth Dimension Technologies' Head/Hand XYZ Tracker (HHT) is a 3 DOF ultrasonic translation tracker system. The system consists of three transmitters, a small receiver unit, a PC interface card, and two interface cables. It is capable of tracking position (x , y , and z) of up to three objects (e.g., head, left hand, and right hand) simultaneously. The tracking system has a worst case accuracy of 20 mm (at 2 m) with an accuracy of 4 mm when the tracked object remains stationary. It has been specifically designed to provide position tracking for the Fifth Dimension DataGlove, see Section 5.1.2.1.

Some specification details for the tracker are provided in Figure 43. The price for the basic HHT Tracker (with one receiver) is \$345, additional receivers cost \$85 each.

Specification	
Update Rate	20 Hz
Accuracy	20 mm worst case (at 2 m), 4 mm for stationary object
Tracking Distance	Up to 3 m

Figure 43. Head/Hand XYZ Tracker

3.1.1.22 GP12-3D (Freepoint 3D)

The GP12-3D, by Science Accessories Corporation, is an ultrasonic product with an update rate of up to 75 Hz divided by the number of emitters being tracked. It goes by the marketing name of

Tracker	Working Volume
Freepoint 3D XL-1	3.25 x 3.25 x 3.25 ft
Freepoint 3D XL-2	8 x 8 x 8 ft
Freepoint 3D XL-D	16 x 8 x 8 ft

Figure 44. Working Volume for Freepoint Trackers

Freepoint 3D and comes in three models. These models all provide a resolution of 0.002 inches, but vary in working volume as shown in Figure 44. Up to 4 emitters may be used together, allowing Freepoint 3D to be used for multiple-unit tracking. The cost is \$4,995 - \$6,995, depending on the model. Each transmitter is only tracked with 3 DOF (x, y, and z).

3.1.1.23 Logitech 3D Mouse and Head Tracker

The Logitech 3D Mouse enables users to provide direct inputs into a VE system. It can operate on a desktop in a similar manner to a traditional mouse, or when raised off the 2-D plane to provide 3-D spatial information. User inputs are specified by means of five buttons on the mouse. In addition the 3D Mouse supports an audio function that allows it to be used as a microphone.

Specification	
Sampling Rate	Up to 50 Hz
Linear Resolution	0.004 in
Angular Resolution	1/10°
Tracking Speed	Up to 30 in/sec
Tracking Space	5 ft, 100° cone

Figure 45. Logitech 3D Mouse

With respect to its tracking function, the 3D Mouse is a low cost ultrasonic device that operates with 6 DOF. The position reference array transmitter is a triangle of three ultrasonic speakers which send signals to a receiver. The receiver itself has a triangular set of three microphones which sample signals from the position reference array. The receiver and transmitter are both connected to a control unit with a CPU that converts the receiver information into position, orientation, and button information. Figure 45 provides further details. An independent study by Ware and Balakrishnan (1994) found the device lag for the Logitech 3D Head Tracker to be 72 msec.

A variant of the Logitech system that may be used as a head-tracker is a small triangular-shaped device that attaches to a HMD. Specification details for the head tracker are the same as those given for the 3D Mouse, with the addition that the tracking space for the head tracker is a linear 5 ft, with a 100° cone. Logitech 3D Mouse and Head Tracker currently costs \$1,599.

3.1.1.24 MotionPak

MotionPak, a Systron-Donner product, is a 6 DOF inertial sensor cluster used for measuring linear and rotational motions. It is also suitable for tracking human motion. Three solid state gyros are used to sense angular rates and servo accelerometers sense linear accelerations. MotionPak weighs 32 ounces. According to Strickland et al (1994), it is both heavy and expensive (\$10,000).

Systron-Donner also market the QRS and GyroChip family of sensor products that employ a pair of micro-machined vibrating quartz tuning forks to sense angular velocity through a deflecting force acting on the body in motion due to the Coriolis principle. The GyroChip weighs 100 grams and costs \$1,000 in quantities of ten to fifty. It has a resolution of 0.004°/second for motion in the 100°/second range, but can handle movement up to 1000°/second. Typical drift after stabilization is reported to be 5°/hour. QRS sensors are used in the MotionPak.

3.1.1.25 GyroPoint Pro

Gyration, Inc.'s GyroPoint Pro is a mid-air mouse that operates in a functional manner similar to a mouse or trackball, but does not need a work surface on which to operate. Unlike line-of-sight IR technology, GyroPoint Pro does not need to be carefully aligned to a receiver. It is compatible with the standard Microsoft or Apple Macintosh Mouse Driver, or Philips CD-i interface. Radio technology provides wireless operation of the GyroPoint Pro within 75 feet of its receiver, even transmitting through windows and walls. GyroPoint Pro cost is \$399.

The GyroPoint Pro uses the GyroEngine sensor developed by Gyration, Inc. This is a miniaturized spinning wheel gyroscope and comes in two models, the Vertical GyroEngine (Model GE9100-A) and the Directional GyroEngine (Model GE9300-C). The Vertical GyroEngine measures roll and pitch whereas the Directional GyroEngine measures heading. Both contain gimbals that are optically encoded, and digital phase quadrature output is available at the GyroEngine 6-pin connector. The Vertical GyroEngine gimbals permit freedom of movement in 360° of roll, measured on the outer gimbal, and a $\pm 80^\circ$ of pitch, measured on the inner gimbal. The Directional GyroEngine gimbals permit freedom of movement in 360° of heading on the outer gimbal, and a $\pm 80^\circ$ on the pitch and roll axes.

The sensor offers an angular accuracy of 0.1° for normal head motion, with the ability to handle head accelerations up to 1000°/second, and a drift rate of 0.5°/minute to 2°/minute. It has a 2 x 1.5 inch size and weighs 2 oz. GyroEngines cost \$295 each.

3.1.2 Current R&D in Head Tracking

This section describes several R&D efforts that are underway. A topic not considered is facial tracking. For those interested, however, a notable R&D effort in this area is underway by the Interactive Systems Laboratory (INTERACT) at Carnegie Mellon University. These researchers are developing a system that tracks human faces with a computer-vision technique based on face color distributions.

3.1.2.1 NASA Ames Research Center

One goal of the research at NASA Ames Research Center, Numerical Aerodynamics Simulation Systems Division, is to study calibration methods for reducing distortions

that result from magnetic trackers. This distortion is significant at distances of greater than 45 to 50 inches, and is sensitive to location since magnetic sensors are very sensitive to metal and electrical devices. The researchers, led by Dr. Steve Bryson of Computer Sciences Corporation, have looked at three methods for reducing static distortions: least-squares polynomial fit calibration, linear lookup calibration, and bump lookup calibration. Of these, 4th order polynomial fit had the best overall behavior, while the bump lookup calibration was superior for tracking very short distances. This early work suggested further study in two directions. First, pursuing the success of the polynomial technique by investigating the use of a 3-D spline calibration (a combination of global polynomials and lookup tables), and study of the weighting and interpolation for the lookup calibration. Second, refinement of the bump lookup calibration method to handle overly distorted data sets.

Current work includes studying dynamic distortion in position data. In this case, the researchers are investigating the use of calibration methods such as Kalman filters. Cross-coupling between these various distortions is also a topic of study.

3.1.2.2 Massachusetts Institute of Technology (MIT), Research Laboratory of Electronics

Dr. Eric Foxlin and Dr. Nat Durlach at MIT, Research Laboratory of Electronics, have developed a prototype inertial navigation system that uses Systron-Donner GyroChips (see Section 3.1.1.24). The goal of their research is to investigate the applications of inertial navigation systems to head tracking to overcome some of the limitations of current trackers. Tracking only orientation, not position, the prototype achieved 1 msec latency, unlimited tracking, 0.008° resolution, and 0.5° absolute accuracy (no drift). The researchers plan further work to develop a second prototype and incorporate it into a complete VE system. They also plan to extend the inertial tracker to 6 DOF tracking.

3.1.2.3 Computer Graphics Systems Development (CGSD) Corporation

A potential solution to the limitations of magnetic trackers is to use a hybrid tracker that exploits the complimentary nature of inertial and magnetic tracking. In this way, the inertial tracking can provide the short term accuracy needed to average out the noise in the magnetic tracker, whereas the magnetic tracker provides the accurate average position needed to eliminate inertial drift.

In a subtask of the Virtual Cockpit project funded by the Simulation Training and Instrumentation Command (STRICOM), Computer Graphics Systems Development (CGSD) Corporation is developing a high accuracy, low latency hybrid tracker for VEs that combines inertial sensing elements such as accelerometers and angular rate sensors with a magnetic tracker. The goal is to develop a 6 DOF tracker with low latency and increased immunity to electromagnetic interference to be used for head, hand, and foot tracking.

A key element of the development approach is the use of Kalman Filtering, from the field of aerospace systems, particularly guidance and navigation systems, to combine the

data. In this way, the data from two sensors can provide better results than could be obtained from each sensor separately. For example, if the angle measurement was slightly in error, gravity would be integrated causing a position drift. The Kalman filter uses the position error derived from comparison to the magnetic tracker to assess the angle error, and correct the error accordingly.

CGSD Corporation has constructed the prototype hardware and is currently developing the software, and expect to have the completed unit ready for demonstration in early 1996. Assuming the tracker meets expectations, it could be commercialized as early as the second half of 1996. Researchers also hope to develop a hybrid inertial/optical tracker that could provide high accuracy over large areas at a cost much less than current optical-only trackers.

3.1.2.4 University of North Carolina

Under the leadership of Dr. Gary Bishop, researchers at the University of North Carolina (UNC) at Chapel Hill are conducting a program of research into wide-area tracking, also called ceiling tracking. The UNC optoelectronic tracking system features LEDs mounted in a ceiling superstructure and upward-looking position sensors, based on lateral-effect photodiodes, mounted on a HMD that a user wears when walking under the ceiling. The system works on the principles of celestial navigation using the fact that the locations of the ceiling's LEDs are known and thus serve as navigation beacons. The geometry of the sensors on the HMD is also known. When the sensors see several LEDs as the HMD wearer moves, a real-time multiprocessor system computes the position and orientation of the user's head. The optical beacon tracking technology means that data are free from the distortions commonly arising in the use of magnetic trackers. These data are then sent to a graphics application which renders the images the user sees in the HMD.

Although the optical tracker used in the wide-area tracking system gives satisfactory accuracy over a large working volume, there are various reasons why its design does not lend itself to hand tracking. One reason is the bulkiness of the cameras. Another arises from the geometry of the situation. For example, the user's body may obscure the hand's "view" of the ceiling, and the hand may not be held upright. Additionally, possible hand movements impose dynamic range requirements on photodiode sensitivity (because of changing distances from the ceiling). Since magnetic trackers usually provide satisfactory performance within a small tracking volume, the wide-area tracking system supports hand tracking by placing a magnetic source on the head mount. Thus, the optical tracker reports the head location in ceiling space and the magnetic tracking system reports the hand's location in source space. Change-of-coordinate transformations among these systems are performed to get the hand's location in ceiling space.

In recent work, the researchers have developed new algorithms for extracting user motion from a sequence of LED sightings. This new method updates the estimate of the user position and orientation on every sighting and has allowed computing the position and

orientation of the headset at greater than 1000 Hz with a delay of less than 2 msec. They have also designed a new head-mounted tracker assembly that is called the "HiBall." This apparatus is about 1 inch in diameter and integrates six cameras with digitization and communication circuitry. When the first prototype is completed, the HiBall will be used to replace the bulky off-the-shelf cameras and the custom electronics in the back pack of the old system. Besides the obvious advantages of size and weight over the old system, the new system will be much faster and more rigid. The HiBall may be useful for hand tracking, as well as head tracking.

There are several other on-going efforts in the wide-area tracking program. In one, the researchers are working to improve the tracking performance of the system from the current resolution of <2 mm in position and 0.2° in orientation, to reach their goal of tracking accurate within 1-2 mm and 0.1° . The researchers also are expanding the existing 16 x 18 foot area ceiling to cover a 16 x 30 foot area. This enlargement is based on a new ceiling panel design that fits in a standard ceiling grid without the need for a heavy metal superstructure. In additional efforts, the researchers are investigating problems for tracking in unenhanced environments (including outdoors) and have ongoing research into the use of inertial sensors in a hybrid configuration with outward looking optical sensors. Other research is investigating tracking systems based on passive targets instead of active LEDs.

3.1.2.5 Artificial Reality

In 1982, Dr. Myron Krueger proposed developing VideoDesk to Defense Advanced Research Projects Agency (DARPA) (now ARPA). In 1987, it was implemented as part of an National Science Foundation (NSF) Small Business Innovation Research (SBIR) project. VideoDesk is currently the focus of another SBIR Phase 1 study that is investigating how to present maps to blind people.

VideoDesk consists of a light table with a camera mounted above it that is aimed down at the user's hands, which rest on the desk's surface. The silhouette image of the hands appears on a monitor, also on the desk. In this way, the user's hands are superimposed on an application and, by means of a gesture interface, he can use the image of a finger to point, draw, or write.

To operate VideoDesk in 3-D, a sample plane is placed anywhere in the volume, in any orientation. Then, the live image of the user's hands is projected onto it, where they can be used to perform 2-D pointing and drawing in the sample plane. By using a second camera, it is possible to perceive the user's hand in 3-D. The most promising applications envisioned for VideoDesk are teleconferencing and teletutoring. ISDN, which enables computers to communicate over phone lines in real-time, is expected to facilitate the development of these applications.

A separate effort, the Project on Biomedical Technology, being sponsored by ARPA, includes the development of a wide-area head tracking system that can track a 30 square foot area. This tracking system combines onboard (relative) tracking and external

sensors. Angular measurements will be achieved by placing sensors on people, whereas relative movement will be tracked using external sensors placed in the environment. The system will use optical sensors that the company is developing from off-the-shelf components. The goal is to reach a data rate of 200 Hz, if possible.

3.1.2.6 Massachusetts Institute of Technology, Media Lab

Using an environment for immersive virtual experience based on computer-vision techniques, the Artificial Life Interactive Video Environment (ALIVE) project at the MIT Media Lab enables interaction between people and agents via natural hand and arm gestures, without the need for HMDs and data gloves. It is under the leadership of Dr. Pattie Maes. The system uses a passive camera tracking system mounted above a projection screen to isolate the image of the user from the background room and locate the user's head, hands, and feet for interaction with the environment. The image of the user, composited with the VE, is flipped horizontally and projected onto the screen, creating a "magic mirror" effect.

Previously, the passive tracking was implemented using a special-purpose vision box by Cognex, which did background subtraction and hand tracking by direct manipulation of the bitmap. Composition of the VE and the real-world room was achieved by chroma keying, which necessarily kept the user in front of the computer graphics. Because of the nature of the ALIVE system, complex heuristics had to be handwritten for the vision system to understand different and unusual positions that people might assume for interaction with the agents. Examples include bending over to the side and squatting down; in both cases the hands are not where the system would generally expect.

The current system uses a digitizer on a Silicon Graphics Indigo2 to reimplement the background subtraction algorithm in software. This allows not only portability but also flexibility of the dependent algorithms. The resultant bitmap from the background subtraction is converted into a polygon which can be rendered into the 3-D VE, allowing occlusion of objects by a user and vice versa. As a result of this conversion process, extremities are essentially found automatically, reducing the hand tracking problem to a problem of classification. This system is not only more general than the previous one, but also drastically reduces the reliance on heuristics.

Researchers are exploring novel applications in the area of training and education, entertainment, and digital assistant interface agents. Current ALIVE worlds include a virtual dog the user can play with and video-game creatures the user can interact with. Another world where a synthetic animated aerobics instructor gives the user detailed personal feedback is under construction.

3.1.2.7 Sony, Computer Science Laboratory

Dr. Jun Rekimoto at Sony Computer Science Laboratory is working on VirtuaHead, a head tracking method using computer vision to support desktop VE. This system employs

a video camera placed on top of a normal CRT monitor to capture user images and track the position of the user's head while he is seated at a desk. The user does not need to wear any special gear. The approach is based on some simplifying assumptions. First, that an approximate position of the user can be assumed to be in front of the screen and calculation of head orientation can be ignored. This allows estimation of the head position to be made in real-time using the two simple image processing techniques of frame subtraction and template matching. Frame subtraction is used to subtract a pre-stored background image from the captured image to detect the user's face area. Then a template of (part of) the user's face is matched against the remaining image to identify the center of the user's face. The resulting u, v position forms the basis for final calculation of head position.

Dr. Rekimoto (1995) reports on an experiment where these researchers looked at how their optical head tracking helped a user's 3-D perception. The VE used for this experiment presented three wire-frame trees positioned at the vertices of an equilateral triangle, one of which had a leaf on a particular branch. The experimental task was to identify which of the trees contained the leaf, comparing conditions of using head tracking or not. The data collected for six subjects showed that while subjects using head tracking took longer to report answers, they had significantly lower error rates. In addition, the researchers noticed that subjects without head tracking often gave up in difficult cases, while subjects with head tracking kept trying by repeatedly moving their heads.

Currently, VirtuaHead does not detect the distance between the screen and the user. Since this distance can change, the tracker can report inaccurate positions. The researchers are working on a solution to this problem. The approach being taken is to estimate the distance based on the size of the face image.

3.1.2.8 Siemens' Central Research and Development

Dr. Christoph Maggioni at Siemens Central Research and Development has developed a computer system called GestureComputer that is able to work in real-time on a standard workstation under noisy and changing environmental conditions, and that detects the 3-D position and orientation of the human hand, as well as the position of the head. The system uses video cameras to observe head and hand movements, relieving users from cables, gloves, helmets, or other encumbrances. Image processing algorithms perform head and hand tracking as follows. A color image is segmented to find patches of human skin by using a fast look-up table approach. The contours of the resulting binary image are traced and a new data structure is generated. This contour data is analyzed and special features are extracted in order to detect head and hand positions. Siemens expects to release GestureComputer as part of some of its products during the next three years.

3.1.2.9 Boeing Information and Support Services, CMU, Honeywell, Inc., and Virtual Vision, Inc.

As part of the ARPA Technology Reinvestment Program (TRP), Boeing Information and Support Services, CMU, Honeywell, Inc., and Virtual Vision, Inc. are using an

optical inside-out videometric tracker for an industrial augmented reality application, touch labor manufacturing. The goal is to develop a position-sensing device that tracks the factory worker's head direction. With this information, the display unit can project the information through a transparent display and onto the work surface, as if it were painted on. Workers having this capability do not have to look through books, so their hands are free to work without interruption. Moreover, the system obviates the need for the expensive marking systems now used in aerospace manufacturing. A benchtop prototype was demonstrated in June 1995.

The project is to assist workers assembling aircraft cable bundles. The worker will wear a belt and headband with a high resolution display and small camera (videometric tracker) through which he will look at the board he is wiring. Graphics will be superimposed on the board, showing him where to put the wires. In order to accurately project graphics onto specific coordinates of a workplace, it is necessary to have the coordinates of the workpiece, the display's virtual screen, the position sensor, and the user's eyes in the same coordinate system. The project requires high accuracy, long-range tracking in a high-noise environment.

In order to compute the position and orientation of a camera mounted on the user's head, fiducials (or markers) are mounted in the work environment and their locations accurately measured. Based on where in the field of view the fiducials appear, the computer can calculate the position and orientation of the camera, and therefore the user's head. For the June 1995 benchtop prototype, Boeing used black paper with white spots as fiducials. For factory use, they plan to mount a bright LED on the user's head pointed towards the work piece, and attach retroreflective targets to known locations on or near the work piece. The camera will need an optical filter that will only pass light of the same frequency as the LED.

Computation involves four steps: capture, correspondence, 3-D reconstruction, and matrix computation. In the capture step, a digital representation of the field of view of the camera is obtained, fiducials from the background are extracted, and the 2-D position of the fiducials accurately computed. The correspondence phase matches each fiducial in the camera's field of view with one of the physical fiducials. Once the location of the fiducials is known on the 2-D screen of the camera, 3-D reconstruction computes the 3-D location of each camera fiducial in the coordinate system of the camera. Finally, given the 3-D locations of the camera fiducials in camera coordinates, and their corresponding 3-D location in real world coordinates, the matrix computation step computes the transformation from one coordinate system to the other. This transformation embodies the position and orientation of the user's head.

Much work on this project remains to be done, both in constructing a production quality unit, and in characterizing the algorithm. Researchers will continue analyzing errors, introducing noise in the camera parameters, lens aberrations, and errors in fiducial placement. Measuring the resulting degradation in accuracy will help them to focus on the major sources of errors. The researchers will also develop methods to calibrate these sour-

es of errors. Two methods that are being explored are auto-calibration and optimization calibration.

3.1.2.10 University of Washington

Recognizing the need for tracking technology which addresses some of the limitations of current head-tracking technology, such as price, accuracy, resolution, and range, the University of Washington, Human Interface Technology Laboratory (HITL), has conducted work on a fast, wide area, multi-participant tracking system. The system used a swept laser fan with multiple sensors mounted on a helmet to track position. The efforts were suspended when funding ran out two years ago, but HITL is now revising the project in hopes of attracting new funding.

HITL researchers are not sure what method of tracking they will now pursue. They are studying the requirements and looking at a number of methods, including the swept laser fan. Advantages of the swept laser fan include the ability to track multiple participants in an area (no real limit). Moreover, no feedback is required between the scanner and the object being tracked. Finally, simplicity of system setup and calibration make it an attractive method. The disadvantages are its limited update rate (60 Hz or so), the power required in the swept beam, and retaining accuracy at the far end of the range. The goal of HITL researchers is to have a system that is accurate to 1 mm over a range of 10 x 10 meters.

3.2 Eye Tracking

Eye-tracking technologies measure the direction the eyes are pointed out of the users' head by detecting the movement of the fovea. This information is then used to determine the direction of the user's gaze and to update the visual display accordingly. General approaches are optical, electroocular, and electromagnetic. The first of these, optical, uses reflections from the eye's surface to determine eye gaze. Electroocular approaches use an electrooculogram (EOG) via skin electrodes that provide measurement of the corneoretinal potential. Finally, electromagnetic approaches determine eye gaze based on measurement of magnetically induced voltage on a coil attached to a lens on the eye.

3.2.1 Commercially Available Eye Trackers

Most commercially available eye trackers are optical and those covered here are products from ISCAN Inc., LC Technologies, Hughes Trainer-Link Corporation, and Forward Optical Technologies, Inc. BioMuse from BioControl Systems, Inc. is the exception and the only product identified that uses the electroocular method.

3.2.1.1 BioMuse

The BioControl Systems Inc. BioMuse System eye controller uses an EOG as the source signal for deriving eye movement information. The EOG itself is derived from the resting potential (known as the corneal-retinal potential) generated within the eyeball by the metabolically active retinal epithelium.

Historically, the EOG has been used as an indicator of eye movement in physiological research studies and in the clinical environment. This relatively common use of the EOG for eye movement evaluation is due to the fact that the technique is noninvasive and the most cost effective and practical of the eye tracking technologies. However, standard amplification and recording techniques present the clinician with several artifacts and technical problems in the evaluation of eye movement for diagnostic purposes. For example, eyelid movement and ocular muscle electromyogram (EMG) contaminate the EOG record. In addition, the recording baseline is unstable due to electrode drift, and repeated calibration is required with adaptation to ambient light. Further, the vertical movement record is unreliable due to eye movement associated with reflex blinking.

For these reasons, traditional EOG recording technology for simple detection of eye movement is unsuitable for an eye controlling system where the user is the initiator of action using eye movements. BioMuse resolves several of the problems by using a DC coupled amplification system to acquire a signal that can be used to indicate the steady state displacement of the eyeball. However, a DC coupled physiological recording system still exhibits an unstable recording baseline caused by electrode drift. To deal with the problems of electrode drift, BioControl Systems, Inc. has developed techniques using fuzzy classification and pattern recognition which are able to greatly reduce this problem.

With proper placement of recording electrodes, vertical and horizontal eye movements can be mapped to move video objects around the screen. The company's 2-D eye controller uses a lightweight headband to position the EOG electrodes. Three electrodes are positioned on the forehead to track horizontal eye movements, and one electrode is positioned below the eye for the vertical channel. For 3-D applications a different headband configuration is required which uses five electrodes on the forehead and two below the eyes, one on each side of the face. This configuration of EOG electrodes is necessary to create two independent horizontal and vertical channels for each eye. With individual measurements for the horizontal and vertical movements, an ocular convergence signal can be derived, and this convergence signal is the basis for the 3-D controller. As the eyes focus on a near field object, they converge, or point inward and, as the object moves into the distance, the eyes diverge slowly until they are parallel at optical infinity. The depth, or third dimension channel, is unique to this patented eye controller system.

The BioMuse System enables an individual to use their nervous system to control virtual objects. The BioMuse product and software tools (libraries) allow users to develop their own motion capture system. The software is in Version 5.0 and provides 3-D depth of field with enhanced drift control to allow greater accuracy. Also included is a Microsoft Windows MIDI (musical instrument digital interface) application that enables a user to play up to 128 musical instruments with obstinacies generated by muscles, eyes, heart, or brain. A video game interface enables users to navigate with their eyes, fire a gun by natural hand motion, and move forward by walking on the spot. The price is \$19,800 for single quantity purchases.

3.2.1.2 Headhunter Head and Eye Tracking System

The ISCAN, Inc.'s Headhunter Head and Eye Tracking System employs helmet mounted eye tracking technology to non-invasively monitor the position of the subject's eye with respect to a miniature imaging sensor mounted on the helmet. This system uses the RK-426 Pupil/Corneal Reflection Tracking System, a real-time digital image processor that simultaneously tracks the center of the pupil and a reflection off of the cornea from the IR light source. RK-426 is a dark pupil tracking system which enables it to perform in high illumination environments with virtually any user. The RK-520 Autocalibration System uses raw eye position data generated by the RK-426 to calculate the user's point of gaze with respect to a scene being viewed and is included in the Headhunter. The tracking system computes eye position at a 60 Hz rate and the eye's point of regard is determined to an accuracy of better than 1° of visual angle over a $\pm 25^\circ$ range. The entire Headhunter System can cost \$30,000 to \$40,000. Cost for the independent RK-426 is \$13,000, or \$7,800 in quantities of five to ten. RK-520 Autocalibration System costs \$6,800. (Other ISCAN eye tracking processors include the RK-406 Pupillometry System, the RK-416 Pupil Tracking System, and the RK-436 Pupil/Dual Corneal Reflection Tracking System).

3.2.1.3 Eyegaze System

LC Technologies' Eyegaze Development System is a workstation for both developing and running custom eye tracking applications. It is a tool for measuring, recording, playing back, and analyzing what a person is doing with his eyes. The Eyegaze System uses the Pupil-Center/Corneal-Reflection method to determine the eye's gaze direction. A video camera located below the computer screen, or below the work space if the computer monitor is not used, continually observes the subject's eye. An IR LED located at the center of the camera lens illuminates the eye, generating the corneal reflection, and causing the bright pupil effect that enhances the camera's image of the pupil. Image processing software identifies and locates the centers of both the pupil and corneal reflection. Trigonometric calculations project the person's gaze point based on the positions of the pupil center and the corneal reflection within the video image. No attachments to the head are required, however the eye must remain within the field of view of the camera.

Specification	
Sampling Rate	30-60 Hz
Gaze Cone Diameter	80° (typical)
Head Motion Tolerance	1.5 (lateral), 1.2 (vertical), 0.5 (long.) in

New eye tracking data are generated each 30th or 60th of a second in synchrony with either the frame or field rate of the video camera. To allow the eye tracking software to run concurrently with an applications program, the eye tracking software runs on an interrupt basis, receiving its interrupt from the camera's frame-grabber board and executing as a stay-resident interrupt service routine. The Eyegaze Development System costs \$21,500. Further details are given in Figure 46.

Figure 46. LC Technologies EyeGaze System

3.2.1.4 Dual-Purkinje-Image (DPI) Eyetracker

The Dual-Purkinje-Image (DPI) Eyetracker was developed at SRI International from 1965 to 1988. In 1988, the technology was licensed by SRI to Fourward Optical Technologies, Inc. which has continued its development and marketing. Currently, Fourward offers two Eyetracker models: the top-of-the-line Generation 5.5 DPI Eyetracker and the DPI Eyetracker 1000. Both are 2-D instruments in that they track only the horizontal and vertical movements of the eye.

The concept behind DPI is as follows. Light rays striking the eye produce four reflections, called Purkinje images, from the front and rear surfaces of the cornea and lens. The first (virtual) and fourth (real) Purkinje images lie in the same focal plane. These two images move similarly under translation, but differentially under rotation. The change in their separation is used to determine eye rotation. Thus, the DPI eye tracker tracks the first and fourth Purkinje image; the latter is dim, so bright illumination of the eye is needed. A photocell captures the reflections and also drives a servo-controlled mirror with an analog signal, avoiding the need for discrete sampling. Figure 47 provides further details for the DPI 5.5 Eyetracker.

The monocular DPI Eyetracker 1000 costs \$49,000, whereas the binocular version costs \$99,800. For the DPI 5.5 Eyetracker, the prices are \$60,000 and \$115,000 for the monocular and binocular versions, respectively. Three accessories that are required with

Specification	
Frequency Response	~500 Hz
Output Delay Signal	0.25 msec
Slew Rate	2000°/sec
Noise Level	20 sec of arc RMS

Figure 47. DPI 5.5 Eyetracker either Eyetracker model are the Headrest Accessory Package (\$1,000), the Model Eye with Control Electronics for maintenance, checkout, and calibration of the Eyetracker (\$2,995), and the Video Eye System turnkey system consisting of an IR camera with monitor and optics that is required for alignment and monitoring (\$2,500 when purchased with an Eyetracker).

In order to track in 3-D, the tracker must know not only the 2-D angular orientation of the line of sight but also the distance from the eye to the point in space on which the eye is fixated, that is, the level of accommodation. The attachment of Fourward's Infrared Optometer enables the measurement of the subject's accommodation. This additional product costs \$29,950.

3.2.2 Current R&D in Eye Tracking

This area is seen as one with a limited application area and there is little active research. It is worth noting, however, that two of the research efforts described below do circumvent one of the biggest limitations of eye tracking technology, that is, its intolerance to head motion.

3.2.2.1 Hughes Training-Link Corporation

The Hughes Training-Link Corporation (formerly CAE-Link) Eye-Slaved Projected Raster Inset (ESPRIT) is a simulator foveal projector eye-tracked display system that uses a bright pupil tracker. It provides a high-resolution (better than 4 arc minutes per line pair) picture in an area of interest mode of operation using a General Electric light valve as display source and a helmet-mounted oculometer from Honeywell. The scene is inset into a background by a high-speed servo that operates at a speed faster than eye saccade velocities (saccade movements can reach 700°/second and have accelerations of up to 50,000°/second). The static line of sight accuracy of the azimuth/elevation servos is better than 1 minute of arc, while the dynamic accuracy is better than 3 minutes of arc at an angular velocity of 700°/second. Latency is about 90 msec, a lag that test subjects found acceptable.

While two of these systems were built for the British, they are very expensive (in the millions of dollars) and were never produced on a commercial scale. Future development plans call for development of a proprietary oculometer by Hughes Training-Link.

3.2.2.2 Interactive Systems Laboratories (INTERACT)

Previously under the direction of Dr. Pomerleau and now under the direction of Dr. Waibe, the Gaze-Tracking Team at the Interactive Systems Laboratories (INTERACT), School of Computer Science, Carnegie Mellon University has developed an artificial neural network-based gaze tracking system that can be customized to individual users. A three layer feed forward network, trained with standard error back propagation, is used to determine the position of a user's gaze from the appearance of the user's eye. Unlike other gaze trackers, that normally require the user to wear head gear or to use a chin rest to ensure head immobility, this system is entirely non-intrusive. In his experiments, Pomerleau has achieved accuracy of 1.5°, while allowing head mobility. Average accuracy is 1.7°. In its current implementation as an input device, the system works at 15 Hz.

The researchers hope to increase the system's accuracy through several additions. When using low resolution images, the pupil and cornea do not provide enough information for the neural net to support accurate gaze tracking. In order to obtain more information from the appearance of the eye, the researchers have used the position of the cornea in the eye-socket. However, this method makes the eye tracker less invariant to head position. One method of addressing this problem is training on multiple head positions and necessitates collecting large amounts of data. In the current system, data collection requires approximately three minutes of the user visually tracking the cursor. In this time, 2000 images of the user's eye paired with the position of the cursor are gathered. If the system were to be invariant to distance from the screen, and relative position with respect to the screen, more training image/gaze location pairs would have to be gathered. An alternate method of maintaining position invariance in the eye tracking system is through the addition of extra input units to represent the head position. Because the camera used in this system has a relatively

wide field of view, the same image from which the user's eye is extracted can be used to extract information about head position.

A potential method to rapidly train the neural network for new users may be to use a multiple network architecture. Several smaller expert networks, under the control of an arbitrating network, can be trained on the eye images of different users. The arbitrating network selects the expert network that is yielding the best response or combines the responses of several expert networks. Preliminary results of this approach have yielded noticeable performance improvements. Alternatively, arbitration could be through the use of metrics which estimate the output reliability. As another attempt to make the same network robust to a variety of people, preliminary experiments have also shown that training a large network with the images of several user's eyes improves the performance for each user.

Currently, the INTERACT Lab is combining the gaze tracking work of Dr. Pomerleau with its face tracking work to extract the face image first and then the focus of the eye based on the image of the face. The goal of this is to eventually extend the ability to identify gaze even for a person moving about the room freely (not only for someone sitting in front of the screen in a relatively stable position as is presently the case). Extraction of the eye image is a problem they are now trying to solve through combination of face tracking and segmentation modules.

3.2.2.3 State University of New York

The Computer Science Department at the State University of New York at Stony Brook has produced an in-house EOG-based eye tracking device from inexpensive off-the-shelf components to detect horizontal and vertical eye movements for human-computer interaction. The system is applicable for both VE systems and video games, as well as for the handicapped. For the latter, the goal is to develop an inexpensive system for use by people able to control only their eye muscles.

The project consists of the development of EOG eye tracking pick-up electrodes, electronics hardware and its fine tuning software, as well as the definition of acknowledgeable eye behavior and the establishment of basic protocols governing on-screen and both 2-D and 3-D object selection and manipulation.

The system itself includes a detecting device adapted to detect the bioelectromagnetic signals generated by eye movements. A first processor receives these signals and generates tokens corresponding to pre-defined eye gestures. These tokens indicate the direction of gaze movement, such as north, south, or north-east, the magnitude of the movement, and the type of eye movement, such as smooth pursuit, saccade, or blink. A second processor receives the tokens and generates command signals based on a token correlation protocol. Thereafter, a user interface responds to these command signals with the appropriate control function. Experiments on the ease of use and accuracy of the system were performed using a 3 x 2 two-level boxed menu driven by eye selections. Subjects were able to make correct

menu selections 73% of the time. However, results improved dramatically (up to 99%) when only four corner squares were looked at, as opposed to the two center squares.

In November, a patent was issued for this eye-tracker. The system is currently in its third prototype and the developers are looking to commercialize it.

3.2.3 Summary and Expectations

To date, low-latency, high accuracy systems for head tracking in unprepared, possibly noisy environment do not exist. Most head trackers achieve a large working volume at the expense of accuracy. These are well recognized problems and, for the relatively mature technology areas of mechanical, magnetic, and acoustic tracking, expected to be the focus of near-term R&D. Improvements are particularly likely in the case of magnetic trackers which are widely used and where the field is dominated by two highly competitive companies (Polhemus and Ascension Technology Corporation). For example, the sourceless, non-inertial trackers described in Section 3.1.1.10, Section 3.1.1.11, and Section 3.1.1.12, are a new type of passive magnetic tracker that offers a cheap alternative to more traditional magnetic trackers. Clymer and Graves (1994) believe they have a strong potential to replace existing inertial, gravity, or mass-based technologies. Because the sensor is non-inertial, the viewer is not subjected to screen "slosh" where head motion has stopped, but the screen keeps moving for a period of time. Also, because it passively monitors its position relative to the earth's magnetic field, the sensor does not need to maintain its alignment relative to any other source. It is, however, difficult to mark movement from one magnetic hemisphere to another, a requirement for full 360° operation.

The most significant improvements in tracker performance, however, are expected to come from the use of hybrid trackers where many of the limitations inherent in a particular technology can be overcome by combining the best features of two of more technologies. Several groups of researchers are doing this. In May 1995, a U.S. patent (Patent No. 5,412,619) was issued for a hybrid triaxial accelerometer/ultrasonic tracking system expected to provide position sampling rates of up to several thousand Hz. Another patent (Patent no. 5,422,715), issued in June 1995 describes a hybrid of an optical localization system and independent tilt and direction sensors. While no commercial hybrid trackers are available as yet, this activity indicates that the next few years most likely will see a growing availability of hybrid trackers using inertial technology. Despite its problems, chiefly a limitation to 3 DOF, inertial tracking provides unlimited range, is fast, and free of interferences. Recent use of silicon micro-machining techniques, that has begun to produce very small inertial sensors, and is leading to an overall reduction in product size and cost, also makes inertial trackers attractive.

The number of research efforts and commercial systems that are using computer-vision with tracking implies that this will continue to be an area of slow growth. Several advances are needed. Most of these systems only deal with 2-D gestures, require complex algorithms, and need significant computing power supported by special-purpose hardware.

None of these problems, however, is insurmountable. The long standing trend toward increasingly powerful hardware at cheaper prices should quickly resolve the last problem, and several groups are working on the development of more powerful and efficient algorithms that can deal with multiple features and users. As reported in Section 3.1.2.4, UNC already has algorithms capable of supporting update rates of 1000 Hz with a delay of less than 2 msec. As the computer vision community continues to make advances in algorithms and hardware, the use of computer vision is likely to become prevalent in tracking.

Wide-area trackers are another area where commercial products are unavailable. However, researchers are demonstrating effective systems capable of both head and hand tracking and these efforts may lead to near-term products. Wide-area tracking is likely to become an increasingly important type of tracking, where the lack of tethering and ability to move freely in room sized areas will make it highly desirable for many different types of VE applications. Current limitations in magnetic, line of sight, and electromagnetic interference are being addressed by various researchers and likely to be resolved in the near future. The wide breadth of possible applications for this type of tracking is likely to encourage continued research and development.

Eye tracking is limited by both the current technology and by the nature of human eye movement. Cost and performance improvements are coming, but only slowly, probably owing to the narrow market this technology serves. Accuracy seems to be the main problem. VE applications need to consider the selection of eye-tracking carefully, since using long gaze or blinking to signal intentions is often an unnatural user interface and limited in its interpretation. The two basic approaches use either biosignals or cameras and, as yet, no one technology seems to be favored above the others. The next few years will likely see ongoing research in both areas. In this timeframe, a small number of new products may come to market, but nothing radically different from the current products.

4. AUDITORY INTERFACES

Historically, discussions of realism for simulation have placed little or no emphasis on auditory requirements. Early work on VEs tended to focus almost exclusively on the visual channel, viewing the VE as a 3-D extension of more 2-D simulations. However, human beings are constantly bombarded with auditory stimuli and, for many applications, eliminating this channel in the development of a VE may be inadvisable.

Auditory stimuli do more than increase the realism of a simulation: In many instances they are essential cues for accurate task performance. For example, Begault (1992) discusses research conducted at NASA Ames Research Center where it was found that pilots had difficulty knowing when they had positively engaged a touch screen "virtual" button without a concomitant auditory cue. When a recording of an aircraft switch engaging or disengaging was coupled with the touch panel, the pilots' preference for the interface was increased. This line of thought can be extrapolated one step further. As in any simulation, in VEs there are some stimuli that should not be used due to their potential danger to the individual. For example, the haptic sensations resulting from crashing an airplane or improperly discharging a high voltage could be fatal to the individual. However, the concomitant auditory cues (for example, the sound of the crash or of the electrical discharge) could be used to simulate the event without any danger to the individual. Indeed, Massimini and Sheridan (1993) have demonstrated the value of auditory cues for substituting for force feedback in various manipulation tasks. In other work, Hendrix (1995) has shown that the addition of spatialized sound significantly increased the reported sense of presence in a VE, although it did not increase the apparent realism of that environment. In an experiment at the Jet Propulsion Laboratory (JPL) Advanced Teleoperation Laboratory, researchers found that auditory feedback, given in addition to visual and kinesthetic feedback, speeded the completion of manipulation tasks (Apostolos et al, 1992).

Although auralization—the 3-D simulation of a complex acoustic field—receives the most attention in discussions of audio interfaces for virtual environments, it would be remiss not to include at least some mention of the field of sonification and its application to virtual environments. Sonification is the audible display of data, such as aids for database or map navigation, or the symbolic representation of error messages or data characteristics. For virtual environment applications having to do with complex multivariate inputs, this technique could greatly aid in data reduction.

One of the most often cited applications of sonification is that of an auditory equivalent to the graphical user interface (GUI). An increasingly experienced problem with GUIs

and their associated icons is clutter of the workspace. With the use of sonification and "earcons" for data representation, this clutter could be significantly reduced. Audio objects and icons, representing data, messages, processes, or resources, could be "placed" in 3-D space around the user, and either automatically signal messages or status changes, or reply to user interrogation.

Although there has not, to date, been the volume of research on sonification that there has been on auralization, some studies have been performed that demonstrate its potential usefulness. For example in a direct empirical comparison of physiological data presented via an auditory display versus a standard visual display, Tecumeh and Kramer (1994) found that subjects responded more rapidly and more accurately to simulated operating room emergencies when using the auditory display. They concluded that for systems where large numbers of variables are causally and temporally interconnected in subtle ways, auditory displays may have a distinct advantage over traditional visual displays. This will become increasingly important as the amount and complexity of data that needs to be processed continues to grow.

An interesting phenomenon that has been reported as part of the virtual audio experience is synesthesia, where another sensory organ is stimulated by the 3-D audio input. For example, the experience has been cited for one high-end 3-D audio product that some auditory inputs are perceived as having tactile properties. For example, when an auditory "soda can" is "opened" next to the listener's ear, not only is appropriate sound perceived, but the concomitant bursting "carbon dioxide bubbles" are felt on the skin. Although no experimental studies seem to have been done to date on this phenomenon, it represents a potential area of interest not only for the study of human perception and the interaction between the senses, but also for the applied area of virtual environments.

These various reasons for the use of sounds, combined with advancements in technology that allow more realistic simulation of real-world auditory inputs, has led to research and development for simulation of the auditory channel for VEs. The current technologies are essentially software-based and regarded as highly proprietary by their developers. Hence, it is impossible to provide detailed information about them.

Before discussing the research and products resulting from the application of these principles in the area of VEs, it is first important to understand some of the mechanics of the human auditory system, including how humans hear and the limitations on the stimuli that they can interpret. A more complete discussion of the anatomy and functioning of the human ear can be found in Scharf and Buus (1986).

4.1 The Human Auditory System

The auditory system has three basic functions: (1) to transmit sound through the outer, middle, and inner ears, (2) to transduce sound waves into neural energy in the inner ear, and (3) to perform neural processing within and transmission through the audio-vestibular nerve and four or five neural levels to the auditory cortex.

The outer ear comprises the pinna (or auricle) and the ear canal (or external meatus). When a sound reaches the outer ear it continues through the air in the ear canal where the pinna concentrates it, increases its amplitude, and reflects the sound at the entrance of the ear canal. This results in the intensity of the sound being changed by as much as ± 10 dB. These effects have been found to be greatest above 2000 Hz (Shaw and Teranishi, 1968). This means that a complex sound containing many high frequencies will vary at the entrance to the ear canal depending on the direction from which it comes. The effects caused by the pinna, therefore, provide cues to the location of a sound source and help to give the impression that a sound source is external to the listener. It should be noted, however, that when auditory stimuli are transmitted via headphones, the sound bypasses the pinna and arrives directly at the ear canal, so that most of the effect of the pinna disappears. It is also worth noting that the difference in propagation delay to the two ears is a source of localization, especially at low frequencies.

After the sound passes through the ear canal, it reaches the middle ear. The primary purpose of this complex system is to match the impedance of the air in the outer ear with that of the fluid in the inner ear. This function of the middle ear prevents sound loss resulting from reflection by the denser cochlear fluid, and allows sound to reach the inner ear with little attenuation. From here, sound vibrations are transmitted through the ossicles (three small bones connected to the tympanic membrane) to the inner ear (or cochlea). This, in turn, causes movement of the cochlear fluid, which causes the basilar membrane to vibrate. This results in bending and activation of the hair-cell receptors lying between the basilar and tectorial membranes, which are innervated by fibers of the auditory nerve. Axons of these fibers enter the central nervous system and synapse in the cochlear nuclei of the medulla, causing the sound waves to be interpreted as audible sound.

There are limits on what sound frequencies are audible to the human ear. These are determined by the acoustical and mechanical properties of the ear canal, ear drum, and the middle ear bones, which set limits on the efficiency with which sounds of various frequencies are converted to mechanical vibrations and transmitted to the cochlea. In humans, the auditory apparatus is most efficient between 1000 and 4000 Hz, with a drop in efficiency as the sound frequency becomes higher or lower. The absolute degree of sensitivity of the human ear is quite remarkable; for example, a movement of the ear drum of less than one-tenth the diameter of a hydrogen atom can result in an auditory sensation. In fact, persons with very good hearing can detect Brownian movement in a soundproofed anechoic chamber (Scharf and Buus, 1986). If the ear were more sensitive than this, random Brownian movements would produce a constant sound and would tend to mask auditory stimuli.

Table 5 presents the international standard threshold values for the minimum audible pressure (MAP) and minimum audible field (MAF). It should be noted, however, that precise values are dependent upon a number of factors, including the type of earphones used in the test and the manner in which the earphones were calibrated. Table 5 gives the values for both the Western Electric 705-A earphone calibrated on a National Bureau of Standards type 9-A coupler and the Telephonics TDH-39 earphone. The values in this table extend

only to 8000 Hz since equipment calibration at higher levels is not reliable. In attempts to overcome this problem, it has been found that for teenagers and young adults the threshold at first slowly rises by 6 to 8 dB, and then jumps another 12 to 14 dB between 14,000 and 16,000 Hz (Stelmachowicz, Gorga, and Cullen, 1982; see also Fausti, Frey, Erickson, Rappaport, Cleary, and Brummett, 1979).

Table 5. Threshold Values in Free Field (MAF) and Earphone Listening (MAP)^a

Sound Pressure Level (dB)						
Frequency (Hz)	ISO Minimum Audible Field (Standard Deviation)		Modified MAF	ISO (W.E. 705-A) Minimum Audible Pressure (Standard Deviation) ^b		ISO (TDH-39) ^c
50	41.7	(6.0)	43.5	—	—	—
120/125	21.4	(5.0)	28.5	45.5	(5.6)	45.0
250	11.2	(4.5)	17.5	24.5	(5.0)	25.5
500	6.0	(4.5)	8.0	11.0	(5.4)	11.5
1,000	4.2	(4.5)	4.2	6.5	(5.4)	7.0
2,000	1.0	(5.0)	1.0	8.5	(5.9)	9.0
4,000	-3.9	(8.0)	-3.9	9.0	(7.6)	9.5
6,000	4.6	(8.5)	4.6	8.0	(7.4)	15.5
8,000	15.3	(8.5)	15.3	9.5	(9.9)	13.0
10,000	16.4	—	16.4	—	—	—
12,000	12.0	—	12.0	—	—	—
15,000	24.1	—	24.1	—	—	—

a. Adapted from (Scharf and Buss, 1986), (ANSI, 1969), (Berger, 1981), (ISO, 1961), (ISO, 1975), (Robinson and Dadson, 1956), and (Wessler, 1968).

b. These values are for the Western Electric 7050A earphone mounted in an MX041/AR cushion and calibrated on a National Bureau of Standards 9-A coupler. The standard deviations are associated with the original threshold determinations.

c. MAP values of the Telephonics TDH-39 earphone mounted in the MX-41/AR cushion. (On the basis of new measurements and review of the literature, Robinson et al (1981) suggest that the MAP values for the TDH-39 earphone are 2.0 to 2.5 dB too high at the 500 to 4000 Hz levels.

The figures given in Table 5 are averages. In fact, thresholds are greatly dependent on age and sex of the listener. Numerous studies have shown that hearing deteriorates with age, particularly for men (for example, Robinson and Sutton, 1979; Moller, 1983). Table 6 summarizes the results from Hinchcliffe (1959), who randomly sampled 400 persons from a rural population of 9000. After rejecting ears with any otological abnormality, a total of 645 ears were tested. Subjects included both women and men between the ages of 18 and 74. In all these cases the men had significantly higher thresholds than the women. Threshold increases with age, more so at high frequencies than at low, and more rapidly after 45 to 54 years. Above 1,000 Hz, males usually have higher thresholds than women in most age groups. At all frequencies, the threshold increases continuously with age, with the greatest loss at frequencies above 2000 Hz. In addition, Table 6 shows that at frequencies from 3,000 to 6,000 Hz, women of all ages have lower thresholds than do men.

It also has been suggested that women have lower thresholds at all ages and frequencies, with this advantage increasing with both age and frequency (Cors, 1963).

Table 6. Threshold as a Function of Subject's Age and Stimulus Frequency

Threshold (dB) re Youngest Age Group												
Frequency (Hz)	Age 8-24 (176)		Age 25-34 (104)		Age 35-44 (93)		Age 45-54 (104)		Age 55-64 (74)		Age 65-74 (94)	
125	4.3		5.5		6.1		9.5		13.1		17.1	
	0.0		1.7		2.6		4.8		8.7		10.1	
	-3.9		-2.1		-1.1		1.8		4.6		6.0	
250	2.9		5.3		5.6		7.6		11.6		16.4	
	0.0		1.0		1.7		3.2		6.5		9.6	
	-3.4		-2.8		-1.5		0.1		2.3		4.5	
500	4.0		4.1		5.7		8.7		12.8		20.8	
	0.0		0.7		1.7		3.9		7.0		9.7	
	-2.7		-2.4		-2.0		0.4		2.4		4.8	
1000	3.7		4.1		6.4		9.5		10.0		24.7	
	0.0		1.0		1.7		4.7		5.6		12.8	
	-3.6		-2.4		-2.1		0.8		1.3		5.2	
2000	4.6		4.8		6.9		10.9		14.9	17.9	26.6	41.1
	0.0		0.4		2.5		5.5		8.7	12.1	14.6	25.1
	-4.2		-3.5		-0.5		1.2		4.6	5.7	9.4	15.6
3000	4.6	8.0	6.9	11.1	10.3	16.5	18.2	29.4	20.2	45.3	40.6	53.1
	0.0	2.1	1.5	5.8	5.5	8.6	9.9	18.2	14.8	31.5	19.8	40.9
	-4.0	-0.9	-2.8	1.1	0.4	3.7	4.8	6.3	8.8	18.7	10.1	30.7
4000	4.3	10.2	8.5	12.9	10.0	19.8	18.9	45.3	26.3	59.6	45.6	59.6
	0.0	3.5	3.8	7.5	5.3	12.6	13.2	22.2	19.4	37.8	22.2	45.5
	-4.3	-1.6	-0.2	2.4	1.7	5.4	6.6	12.4	8.7	30.7	12.1	29.9
6000	5.8	9.6	8.8	12.9	13.6	20.9	22.4	36.7	28.7	59.7	47.2	66.5
	0.0	2.5	3.6	5.3	6.2	12.4	11.2	23.1	22.3	49.5	33.9	50.9
	-6.6	-2.6	0.5	-1.2	0.6	5.9	3.5	14.0	11.3	30.6	17.4	34.9
8000	5.6		9.6		15.7		28.4	45.1	39.7	67.8	52.3	68.5
	0.0		3.3		7.2		8.2	20.7	24.7	53.5	42.2	57.2
	-6.3		-5.4		-2.3		2.6	10.7	11.0	30.1	32.2	48.7
12000	9.2		17.6		82.5		58.0		70.0		70.0	
	0.0		5.2		14.4		41.7		64.2		70.0	
	-6.6		-4.0		3.9		19.0		54.2		63.1	

Note: The middle value in each triplet is the median value; the bottom value is the 25th percentile and the top value is the 75th percentile. The number of ears within each age group is given in parentheses. Where two sets of triplets appear together, the set on the left is for women and that on the right for men. The table displays the medians and the 25th and 75th percentiles of the hearing levels relative to the thresholds of the youngest age group at the corresponding frequency. Where single scores are given, there were no statistically significant gender differences at that frequency.

Reprinted by permission of John Wiley & Sons, Inc.

Typically, sounds are described using three variables: pitch (frequency), tone color (spectral content), and loudness (intensity). When one is trying to synthesize sounds within a virtual environment, a fourth variable comes into play—spatial location. This variable is dependent on a number of factors. First, interaural intensity and time difference cues are

essential to determination of auditory localization. For example, when a sound occurs at 45° to the listener's right at 45° azimuth and 0° elevation, the sound is louder in the right ear than in the left. As the intensity differences vary between the two ears, the listener interprets changes in the sound location. In addition, since the path to the right ear is shorter than that to the left ear in this example, it will reach the right ear fractionally sooner than the left ear.

The original theory of sound localization was based on two types of sound measurements. Since wavelengths smaller than the human head create an intensity loss or head shadow at the ear farthest from the sound source, researchers thought that the brain used interaural intensity differences (IIDs) to localize high frequency sound. Interaural time differences (ITDs), on the other hand, were thought to be important for low frequencies since the interaural delay relationships were clear for wavelengths larger than the human head.

This duplex theory, however, does not account for the ability to localize sounds on a vertical median plane with minimal interaural cues. In addition, the duplex theory does not account for the fact that sounds often appear to be coming from inside the head when heard over earphones, even though the appropriate IIDs and ITDs are present. It is now thought that direction-dependent filtering resulting from sound interacting with the outer ears at least partially explains these deficiencies. Research has shown that the pinnae shape the sound waves in highly direction-dependent ways, and are at least partially responsible for the perception that sounds are occurring outside the head.

The synthesis of a 3-D auditory display typically involves the digital generation of stimuli using location-dependent filters. These filters are constructed from acoustical measurements made using small microphones placed in the ear canals of individual subjects. These ear-dependent filters are usually referred to as head-related transfer functions (HRTFs), and act much like graphic equalizers. The HRTFs capture the IIDs, ITDs, and spectral coloration produced by a sound's interaction with the pinnae that are essential for localizing sounds. An alternative method of measuring HRTFs involves the use of a geometric model of the "average" human head, shoulders, and upper torso. This model purportedly yields results comparable to those obtained from sampled HRTFs as discussed above. The modeling process has also been used to construct artificial heads that can be fitted with miniature internal microphones so that live 3-D sounds can accurately be recorded.

However, 3-D effects are lateralized or externalized rather than truly being localized. This occurs for several reasons: First, there is a lack of input from head motion cues. In trying to localize sounds, humans tend to move their heads toward the hypothesized source, thereby increasing the data they receive and interpret concerning spatial localization. Second, modification of the sound wave by the pinnae acts to emphasize some frequency regions and attenuate others. As discussed by Begault (1992), although two sound sources—one at right 60° azimuth, 0° elevation and another at the mirror image position of 120° azimuth, 0° elevation—have the same overall interaural time and intensity difference, the rearward sound has a relatively "duller" quality. In fact, for a given broadband sound

source, each elevation and azimuth position relative to the pinnae has a unique set of spectral modifications. This is due to the complex construction of the outer ears, which impose a set of minute delays that collectively translate into a particular binarual HRTF for each sound source position.

4.2 Commercially Available 3-D Audio Products

A wide continuum of commercial products are available for the development of 3-D sound. These range from low-cost, PC-based, plug-in technologies that provide limited 3-D capabilities to professional quality, service-only technologies that provide true surround audio capabilities. The characteristics of the products identified here are summarized in Table 7, and the products themselves are discussed in alphabetical order below.

In addition to these products, Focal Point 3D Audio license their family of technologies for creating positional sound. The first of the technologies, Focal Point Type 1, uses DSP-based real-time binaural convolution and supports head tracking. Focal Point Type 2 also binaurally positions multiple sounds, but using only software on any standard PC (Focal Point 3D Audio has applied for a patent on this technology). The final technology, Focal Point Type 3, also called the Focal Point Audio Animator, is software-based and automatically creates positional 3-D audio to match objects and motions in 3-D graphic animation.

4.2.1 Acoustetron II

The Acoustetron II, from Crystal River Engineering, Inc., is a stand-alone, turn-key sound system for developing interactive, 3-D sound for real-time graphics workstations. The system is capable of the full spectrum of 3-D sound, including Doppler shifts, spatialization, and acoustic ray tracing of rooms and environments. Accordingly, it can create sounds that originate from exact positions in 3-D space and exhibit Doppler shifts as they travel past a listener. Additionally, sounds exist in a custom acoustic environment—such as a room, a cathedral, or even outdoors—where they bounce off surfaces, travel in the atmosphere, or pass through materials, and are reproduced in real time.

	Specification
Concurrent Wave Files and Spatial Sources per Card	9 at 44,100 Hz sample rate 16 at 22,050 Hz sample rate
Update Rate	44 Hz
Input	64x oversampled, 16-bit A/D converters
Output	8x oversampled, interpolating filters
Stereo Crosstalk	100 Hz-100 dBv, 1 kHz-80dBv 10kHz-60 dBv
Interface	RS-232

Figure 48. Acoustetron II

be presented over headphones, earphones, or speakers. The ANSI C functions of the server allow fast, high-level development of 3-D sound spaces and easy integration into existing

The Acoustron II is controlled from a central simulation 486DX2-based computer over a communication line. The client sends information such as audio source and listener positions to the server via RS-232. The server continually computes source, listener, and surface relations and velocities, and renders up to 16 separate spatialized sound sources accordingly. The audio output can

Table 7. Characteristics of Commercially Available Auditory Products

Vendor	Product	Description	Product Type	Typical Applications	Price
Crystal River Engineering, Inc.	Acoustetron II	An 8/16 channel, turn-key AudioReality Server for use with SGI, Sun, HP, or PC client systems. Full spectrum 3-D sound, including Doppler shifts, spatialization, and acoustic ray tracing of rooms and environments	Stand-alone system	Computer simulation applications (e.g., vehicle and industrial training, location-based entertainment, video conferencing, and architectural walkthroughs)	Base System \$9,995, Development System \$12,495
	Proton	AudioReality plug-in for Digi Design's ProTools TDM system	Plug-in	Creation of spatialized 3-D tracks in the Digidesign ProTools dynamic mixing environment	\$995
QSound Labs	QSystem I	Underlying QSound technology	Algorithm	Video games, PC applications, consumer electronics, and multimedia	Licensable
	QSystem II	8 channel real-time audio processor for professional recording environment	Stand-alone system		Digital System \$8,500, Digital Plus System \$9,750, Digital Plus/Analog System \$17,000
	QCreator	Windows-based audio authoring tool for creation of preprocessed and run-time effects	Authoring tool		Included in licensed package
	QExpander	Plug-in for Sound Designer	Plug-in		\$295
	QSYS/TDM	4 channel software plug-in version of QSystem for use with Mac-based ProTools III	Software plug-in		\$995
	QMixer 32 and QMixer 95	32-bit DLLs for Windows 3.1 and Windows 95 for mixing and interactive processing of multiple WAV files	DLL		Included in licensed package

Table 7. Characteristics of Commercially Available Auditory Products

Vendor	Product	Description	Product Type	Typical Applications	Price
Roland Corporation US	RSS-10	Two channel sound space processor utilizing a new DSP 3-D sound processor. Capable of generating a complete 360° reverb soundscape, including digital processing of reflections, delays, and Doppler effects	Dedicated, single rackmount unit	Professional studios and sophisticated project recording, as well as audio post production, broadcast, and sound design	\$9,750
	SDX-330	Dimensional Expander to produce modulation effects processing	Dedicated, single rackmount unit	Professional recording, public address, and sound reinforcement, enhancement of individual instruments, and personal recording	\$8,500
	SRV-330	Dimensional space reverb unit that provides reverb sounds with a total of 22 reverb algorithms, including sync, mono, and 3-D	Dedicated, single rackmount digital effects unit	Studio recording, public address, sound reinforcement, home recording, and live performance	\$1,295
	SDE-330	Dimensional space delay unit with a total of 19 delay algorithms. Delay effects can have times as long as 2.9 sec and can be combined for a variety of delay effects	Dedicated, single rackmount digital effects unit	Studio recording, public address, sound reinforcement, home recording, and live performance	\$1,295
	SoundStorm 3D	System generates from five 3-D up to thirty-two 2-D simultaneous sounds using Focal Point rendering algorithms. It can be networked for integration into simulations and supports both DIS and SIMNET protocols.	Stand-alone system	Military simulations	\$20,000
Reality By Design					
Audio Cybernetics	VAPS	Process of actively or interactively encoding audio in 3-D	Currently offered only as a service	High-end recording and mixing of 3-D sound	Contact vendor

VEs. Specification details for the Acoustetron II are given in Figure 48. The base system is available for \$9,995.

The acoustic ray tracing mode (4 concurrent wave files and spatial sources per system at 44,100 Hz sample rate) is supported as an option. Wave file recording and editing studio software, high-speed communication protocols, and quad speaker output are also available as options. In addition, there is a development option to control additional input and output devices from the host computer. Finally, it is worth noting that the Acoustetron II is supported by Coryphaeus' EasyScene, Paradigm Simulation's Vega, Sense8's World-ToolKit, Division's dVs, and Autodesk's CDK world building toolkits.

4.2.2 Protron

Crystal River Engineering, Inc.'s Protron enables sound designers to interactively place and move audio sources in a 3-D custom acoustic space. Using Protron, audio designers can create fully spatialized 3-D audio tracks in the Digidesign Pro Tools dynamic mixing environment.

Protron takes monophonic input and adds the psychoacoustic cues and environmental effects that make it appear to be located at a specific point in space, in a specific environment. Simple mouse operations are used to interactively place the source in 3-D; pop-up menus and sliders are provided to customize the acoustic space. Protron features:

- AudioReality sound field synthesis,
- Complete, interactive 3-D source position control,
- User selectable environment size and materials,
- Continuously adjustable parameters,
- Monophonic compatible output, and
- True RMS level meters.

Protron is fully compatible with Pro Tools III, and Pro Tools II with the TDM option. It has the same minimum Mac hardware and system software requirements as Pro Tools. The maximum number of sources which may be simultaneously rendered is equal to the number of available Pro Tools TDM DSPs. The price for Protron is \$995.

4.2.3 Q Products

QSound Labs is an audio technology company that specializes in low cost sound localization and enhancement. Its products range from analog and digital hardware to stand-alone and add-on software. The three major products are QSystem I, a stereo speaker-based sound localization process, QSystem II, a headphone-based sound localization process, and QXpander, a stereo enhancement process for both speakers (QXpander I) and headphones (QXpander II) that can operate on existing stereo material.

The QSystem I process is the fundamental QSound technology. Using multiple monaural inputs, the process produces a stereo output signal that allows each input to be placed anywhere within an 180° arc around and in front of the listener. Although the QXpander is based on the same underlying technology as the QSystem I, it filters existing stereo images to expand the depth and separation of the sound field. QXpander also provides the ability to create 3-D sound in situations where the elements of the audio mix are not individually accessible.

In addition to these basic audio processing technologies, QSound also has available a number of tools. These include:

1. QCreator: Windows-based audio authoring tool for creation of preprocessed and run-time effects.
2. QSystem II: Eight-channel real-time audio processor for the professional recording environment.
3. QSYS/TDM: Four-channel software plug-in version of the QSystem for use with the Mac-based Pro Tools III from Digidesign.
4. QTOOLS/SF: A low-cost plug in for Sound Forge, by Sonic Foundry, that offers QXpander, static QSystem I, and a sample rate conversion tool.
5. QMixer 32 and QMixer 95: QMixer 32 is a 32-bit DLL for Windows 3.1 (with Win32s extensions) for mixing and interactive processing of multiple sound files. QMixer 95 provides the same technologies for Windows 95.

QSystem I is available for licensing and includes QCreator and QMixer, contact QSound Labs for price information. Prices for QSound II start at \$8,500. QExpander and QSYS/TDM cost, respectively, \$295 and \$995.

4.2.4 RSS-10 Sound Space Processor

Preliminary Specification	
Nominal Input Level	XLR +4 dBm (Headroom: 20 dB), Phone -10 dBm (Headroom: 20 dB)
Nominal Output Level	XLR +4 dBm (Headroom: 20 dB), Phone -10 dBm (Headroom: 20 dB)
Source Distance	Max. 81 m (1 cm step), Max. 655 m (8 cm step)
Reverb Time	0.1-40 sec
Room Size	1-100 m
Interface	RS-422A/232C

Figure 49. RSS-10 Sound Space Processor

Roland Corporation US, a wholly-owned subsidiary of Roland Corporation Japan, focuses on high-end recording, sound reinforcement, and broadcasting applications. Products include digital hard disk recorders, 3-D sound processors, and other digital signal processors. The organization uses HRTF sound processing to create 3-D

audio for playback via speakers. Provided the listener's position relative to the speakers is fixed in accordance with Roland's instructions, their transaural crosstalk cancellation techniques address the problem of sound designated for one ear entering the other, thereby producing strong 3-D effects.

The RSS-10 Roland Sound Space Processor is a two-channel system that utilizes a fast, new DSP 3-D sound processor. A complete 360° reverb soundscape can be generated, including digital processing of reflections, delays, and Doppler effect. Using the RSS-10, sounds can be placed or moved above, towards, or around the listener using standard stereo playback. As the sound source moves through the 3-D plane, the reverb reflections move accordingly, in real time. This creates 3-D sound with natural room ambiance. Specification details for the RSS-10 are given in Figure 49. Software control is available for both the Windows and Mac platforms. The price for the RSS-10 is \$9,750.

4.2.5 SDX-330 Dimensional Expander

Another Roland Corporation US sound processor is the SDX-330 Dimensional Expander. This product is designed to produce high-quality modulation effects. It features Roland's proprietary 3-D spatial simulator sound localization technology, and creates unique 3-D effects with conventional 2-channel playback. The high performance capabilities of the SDX-330 come from the newly developed custom DSP chips that perform over 33 million computations per second. This level of digital processing results in enhanced resolution for precise, smooth, and warm effects, performed in a dedicated 384 kBytes of memory. The SDX-330 has discrete stereo processing, with two independent pairs of inputs and outputs that process the left and right channels independently to maintain true stereo localization of the direct sound within the effected sound.

	Specification
Signal Processing	A/D conversion: 16-bit, Delta-Sigma modulation, D/A Conversion: 16-bit, Delta-Sigma modulation
Sampling Frequency	44.1 kHz
Frequency Response	5 Hz to 70 kHz: -3/+0.3 dB (direct), 20 Hz to 20 kHz: -3/+0.3 dB (effect)
Nominal Input Level	-20/+4 dBm (selectable with Input Level Switch)
Input Impedance	300 k Ω (Input Level Switch: -20 dBm), 10 k Ω (Input Level Switch: +4 dBm)
Nominal Output Level	-20/+4 dB (selectable with Output Level Switch)
Output Impedance	1.5 k Ω (Output Level Switch: -20 dBm), 9 k Ω (Output Level Switch: +4 dBm)
Total Harmonic Distortion	< 0.012% (Direct, 1 kHz at nominal output level)
Dynamic Range	100 dB or greater (Direct), 90 dB or greater (Effect)

Figure 50. SDX-330 Dimensional Expander

from a MIDI keyboard by using performance information such as note number, aftertouch, velocity, and control change messages sent from the modulation lever, data entry slider, or (optional) pedals. User patches can be bulk dumped to external MIDI devices, such as a sequencer or personal computer, for storage. Effects patches can be selected and parameters can also be changed in real time from a sequencer, enabling sequencer-controlled effects processing. Specification details for the SSE-330 are given in Figure 50. The cost of this product is \$8,500.

The SDX-330 features 16 different algorithms, including stereo chorus, multiband choruses, classic chorus simulations, rotary, stereo 3-D chorus, and 3-D panner. Additionally, the SDX-330 offers extensive MIDI control. Effects parameters can be controlled in real time

4.2.6 SRV-330 Dimensional Space Reverb and SDE-330 Dimensional Space Delay

The SRV-330 Dimensional Space Reverb and SDE-330 Dimensional Space Delay units are two additional Roland Corporation USA products, that also rely on newly developed custom DSP chips for high speed digital processing. They provide high resolution, and permit the creation of a wide range of effects. Their sizable internal audio memory allows original sound quality to be maintained even after effects processing. Moreover, A/D/A conversion is 16-bit, with 30-bit internal signal processing and a sampling rate of 44.1 kHz. These technical characteristics deliver high-quality effects typically found only in professional recording equipment, including a flat frequency response of 20 Hz to 20 kHz, dynamic range of 90 dB or higher, and signal to noise ratios of 78 dB or greater.

Both units have stereo configurations (equipped with two inputs and two outputs), and can accommodate any input source. Since stereo algorithms perform internal signal processing for left and right channels independently, the exact stereo sound image localization of the direct sound is maintained. The inputs and outputs accommodate professional +4 dBm line level signals as well as the multipurpose -20 dBm level to meet a wide range of applications.

The SRV-330 Dimensional Space Reverb provides reverb sounds, with a total of 22 specially-developed reverb algorithms. These include:

- Sync Reverb: A stereo reverberation algorithm that creates basic reverb types like Hall, Room, and Plate, as well as extra-high density reverb that adds chorus. The unit also includes internal signal processing for both left and right channels.
- Mono Reverb: This algorithm combines two separate reverb sections, such as Hall and Room, to recreate sound environments with complex reverberation characteristics.
- 3-D Reverbs: These unique algorithms are based on Roland's proprietary 3-D spatial simulator sound localization technology. For example, the 3-D ambiance algorithm provides 24 early reflections that can be positioned at 12 locations for high-density reverb effects. The 3-D reverb algorithm adds dense rear reverberation to 12 early reflections positioned at six locations. Based on these algorithms, 300 preset patches are available. Any of these presets can be edited as desired and 100 customized patches can be stored to SRV-300 memory for instant recall. A built-in 3-band parametric pre-EQ permits precise tonal shaping to match the original sound texture and reverb type selected.

Conventional stereo reverberation units localize early reflections only within a two-channel stereo sound field. In real life 3-D sound environments, however, early reflections are localized at various points. Accurate early reflection reproduction recreates realistic acoustic spaces under different conditions. Algorithms that use the 3-D spatial simulator allow the SRV-330 to generate some extraordinary effects. For example, the 3-D ambiance algorithm generates up to 24 early reflections and positions them at a maximum of 12 loca-

tions in a 3-D sound field. The optimal delay time, phase differences, and filtering for each reflection are automatically calculated. Using the SRV-330, it is possible to get a brand-new palette of creative effects with just a conventional two-channel system, including simulation of high-ceiling rooms, hard-walled rooms, and many other specialized acoustic environments.

	Specification
Signal Processing	A/D conversion: 16-bit, Delta-Sigma modulation, D/A conversion: 16-bit, Delta-Sigma modulation
Sampling Frequency	44.1 kHz
Frequency Response	5 Hz to 70 kHz: -3/+0.3dB (direct), 20 Hz to 20 kHz: -3/+0.3 dB (effect)
Nominal Input Level	-20/+4 dBm (selectable with Input Level Switch)
Input Impedance	300 k Ω (Input Level Switch: -20 dBm), 10 k Ω (Input Level Switch: +4 dBm)
Nominal Output Level	-20/+4 dB (selectable with Output Level Switch)
Output Impedance	1.5 k Ω (Output Level Switch: -20 dBm), 9 k Ω (Output Level Switch: +4 dBm)
Total Harmonic Distortion	<0.012% (Direct, 1 kHz at nominal output level), <0.02% (Effect, 1 kHz at nominal output level)
Dynamic Range	100 dB or greater (Direct), 90 dB or greater (Effect)

Figure 51. SDE-330 Dimension Space Delay

ventional 2-channel stereo systems. The SDE-330 has a total of 19 delay algorithms, and includes new algorithms specifically designed for the unit. Delay effects can have times as long as 2.9 seconds and can be combined for an variety of delay effects:

- **Stereo Delay:** Allows realistic sound effects—from simple delays, all the way up to complex cross feedback—to be easily created through separate processing of the left and right channels.
- **Quad Delay:** The SDE-330 includes four delay sections connected in series, to produce dense delays.
- **Multitap Delay:** This algorithm allows delay effects to be maximized, and can use up to eight taps.
- **Gate/Duck Delay:** The built-in gating function allows the creation of special delay effects. For instance, a delay signal can be switched on and off alternately at the gate threshold setting.
- **Pitch Shift Delay:** The SDE-330 allows delay effects to be enhanced through the pitch-shifting of the four delay taps by ± 1 octaves.
- **3-D Delays:** The SDE-330 is distinguished by several algorithms for radical effects. For instance, the multitap space delay allows delay taps to be positioned at different points around a 360° circle of sound. Using the 3-D chorus, richer and fatter textures than those of conventional chorus units can be achieved. This algorithm positions multiple pitch-shifted signals at many different points all around the direct signal.

The SDE-330 Dimensional Space Delay's unique effects algorithms also use Roland's 3-D spatial simulator. For example, the multitap space delay algorithm generates up to eight delay taps that can be positioned at any point in a 3-D sound field. This allows multidelays to echo around a 360° sound field with only con-

Based on these algorithms, the SDE-330 incorporates 100 preset effects patches. More creative effects can be obtained by customizing the effects patches and storing up to 200 user patches in SDE-330 internal memory. The unit also includes a built-in 3-band parametric EQ that provides the versatility to make fine tonal shading or radical effects without having to employ a separate EQ unit.

Both the SRV-330 and SDE-330 are available for \$1,295 each.

4.2.7 SoundStorm 3D

From Reality By Design, SoundStorm 3D is a stand-alone Pentium PC-based system capable of generating from five 3-D up to thirty-two 2-D simultaneous sounds. The system operates in a Unix environment and can be networked with a variety of computers, such as Sun and Silicon Graphics workstations, for integration into simulations. The system hardware platform is augmented with dual 16-bit sound cards, four magnetically shielded speakers, and Ethernet IP networking. It supports both the Distributed Interactive Simulation (DIS) and SIMNET protocols, and provides a sound effect library suitable for military applications. A set of customization tools allows users to record and manipulate additional sounds. The sound generation is performed independently of the simulation and linked to it by associating particular sounds with specific simulation entities. The actual sound generation employs Focal Point binaural rendering algorithms. Sound Storm 3D is available for \$20,000.

4.2.8 Virtual Audio Processing System

The Virtual Audio Processing System (VAPS) of Audio Cybernetics creates sounds containing significant psychoacoustic information to fool human sensory organs into perceiving that the sound is actually occurring in the physical reality of 3-D space. The sounds can be reproduced on a conventional stereo system, surround system, or headphones, without special decoding equipment.

One of the most significant features of VAPS is that it not only requires no special equipment to decode, but also eliminates the sharp "sweet spot" limitations present in many other systems. This results in a wide listening area that allows for more listener mobility than is available with more conventional systems. Since the process is encoded directly onto the recording medium, the need for special processing devices for playback is eliminated. Rather than being based on anatomical simulation (artificial head architecture), the VAPS uses a highly-detailed mathematical model to define the variables of human audition in the creation of 3-D audio effects. The cited results are sounds that reportedly not only are perceived to occur in 3-D space, but that may also have the synesthetic properties that fool the listener into believing that subtle "tactile" perceptions were received.

The VAPS also allows the user to give some spatial qualities to pre-recorded material, or to "move things around in the mix." Additionally, it allows the creation of sophisticated room simulations with the proximity effects that are necessary for VEs.

The VAPS is part of the on-going research of Audio Cybernetics. Other areas being investigated are the development of a 3-D sound recording chip and the feasibility of producing VLSI components for 3-D audio processing. Currently, the VAPS process is only offered as a service. Audio Cybernetics will provide the virtual audio equipment at a daily rental along with a trained engineer. Expert consultation and full virtual audio production services are also available. Contact Audio Cybernetics for pricing information.

4.3 Current Research and Development

Research and development efforts for the generation of realistic 3-D sound have been conducted since the 1880s. Within the past twenty years, these efforts have escalated, primarily due to impetus received from the entertainment industry. With the relatively recent advances in signal processing technology, acceptable results in 3-D sound are now available for reasonable prices. As a result, increasing numbers of potential applications are being found for this technology. However, although the current state-of-the-art in virtual audio technology is far advanced from where it was even ten years ago, many facets still are not well understood, both in the area of basic research concerning human auditory perception and in the way technology can be improved and applied. The following paragraphs discuss some of the current work in the field of virtual auditory displays that is being performed today. Although other work is also continuing, no information was available at the time of this writing.

4.3.1 NASA Ames

A pioneer in the field of virtual audio, NASA Ames has been working with both the University of Wisconsin - Madison and Crystal River Engineering to develop increasingly sophisticated audio spatial displays, and use these in practical, real-world applications. As a result of their efforts, the Convolvotron, the world's first real-time, 3-D acoustic display, was developed in 1987. Recent work has included investigations of heads-up auditory displays for traffic collision avoidance systems (Begault, 1993), localization in virtual acoustic displays (Wenzel, 1992), multi-channel spatial auditory displays for speech communications (Begault and Erbe, 1993), localization using non-individualized HRTFs (Wenzel et al, 1993), virtual acoustic displays for teleconferencing (Begault, 1995), and headphone localization of speech (Begault and Wenzel, 1993).

Currently, NASA Ames is performing multiple projects relating to the development of virtual acoustic environments, and has recently been awarded a contract to perform both basic and applied research and technology development to implement 3-D auditory displays for improved operative efficiency and safety. As part of this program, NASA Ames will conduct perceptual studies of human sound localization using techniques developed for real-time synthesis of 3-D sound over headphones and apply this knowledge for both enhancing and perceptually validating the advanced acoustic display systems that have been developed as part of their ongoing spatial sound project.

The binaural listening system will enable an astronaut, ground-controller, or other human operator to take advantage of their natural ability to localize sounds in 3-D space, and is intended to be used to enhance situational awareness, improve segregation of multiple audio signals through selective attention, and provide a means of detecting a desired signal against noise for enhanced speech intelligibility. NASA Ames also plans on developing an in-house capability to measure HRTFs and a real-time room modeling technology in the near future.

4.3.2 Naval Postgraduate School

As part of their Naval Postgraduate School Networked Vehicle Simulator (NPS-NET) research effort, the Naval Postgraduate School is another institution currently investigating the practical application of virtual audio technology. Overall, the research effort is addressing many issues, including large-scale networking of VEs, representation of the human body in VEs, and the integration of hypermedia into VEs. The NPSNET itself was developed as student-written, real-time, networked software running on commercial, off-the-shelf workstations. Although originally envisioned as a low cost, government-owned, workstation-based visual simulator, it has evolved to include many facets of VEs, including virtual audio.

The NPSNET Polyphonic Audio Spatializer (NPSNET-PAS) can play sounds either in the spatialized sound mode or directly from the computer's built-in sound board. It keeps track of each sound occurrence in the virtual world (for example, a detonation or vehicle) by "listening" to the packets on the network. The program then determines whether the source of the sound is within hearing range of the player and calculates the correct volume and direction of the sound to be played by the Emax II. The sound is delayed to simulate the correct distance between the occurrence and the listener.

The previous version of the MIDI-based sound system for NPSNET could only generate aural cues via free-field format in 2-D. To increase the effectiveness of the auditory channel, a sound system was needed that could generate aural cues via free-field format in 3-D. To do this, hardware limitations of the NPSNET-PAS sound generating equipment were identified and more capable off-the-shelf sound equipment was procured. In software, new algorithms were developed to properly distribute the total volume of a virtual sound source to a cube-linked configuration of eight loudspeakers and to enhance the ability to localize a sound source. Synthetic reverberation using digital signal processors was added to enhance perceptual distance of the generated aural cues.

4.4 Summary and Expectations

Virtual audio has a value far beyond the music industry for which it was originally developed or for gaming where it is commonly used. Many other applications have been found in such fields as medicine, training and simulation, and architecture. Further, as the technology matures, it is likely even more applications will be found and it will become widely used in VEs.

Although high-end virtual audio approximates the real world, the technology is still far from perfect. In the near future, work is expected to continue on improving the realism and full-surround capabilities of the technology. To do this, better algorithms need to be developed, based on a more thorough understanding of how humans perceive sounds. Other research that needs to be accomplished is a determination of what auditory stimuli are necessary to simulate various environmental sounds. This is closely tied to the need to determine how "realistic" an auditory simulation must be in order to result in the desired effect. In addition, generic HRTFs that maximize accurate localization of sounds in space need to be developed and made publicly available. Linked to this is the need to develop refined algorithms for use in recording sounds, and mathematical models to simulate human hearing organs. Other issues that must be addressed include better control of ambient noise that distracts from the reality of the virtual environment, elimination of unwanted reflections, and technology that will allow the listener to move in relation to the sound source without noticeable distortion of the sound quality.

Since the digital synthesizers used in virtual audio were originally developed for the music industry, synthesized speech and sound effects are not well developed. Although there has been much research in this area, this is a difficult problem and further work is required before automatic virtual audio can be produced in real time. This work needs to address the necessary spectrum of sounds and how they are affected by changes in stimuli, and algorithms need to be created. Similarly, much work still needs to be done in the area of speech synthesis. The significant factors in natural speech yet have to be identified, and ways to synthetically reproduce natural speech without significant deterioration in perceptual quality developed.

Currently, the use of high-end virtual audio technology requires specialized technicians. As the technology matures, one of the expected trends will be an increase in the ease of use. This will include more user-friendly interfaces for the technology, such as standardized options that allow inexperienced or non-professional users to easily approximate professional results, and better human-machine interfaces for power users and professionals who need or desire a wider range of options. As part of this trend, generic HRTFs will need to be synthesized, or current HRTFs modified, so that they can be easily applied by inexperienced users.

Finally, validation studies need to be performed to determine the utility of virtual audio in various applications. It is only in this way that the technology can be fine-tuned and properly applied.

In the near future, as digital signal processing technology becomes less expensive, it is expected that virtual audio will become more widely available at a lower cost. This is happening to some extent already with many dedicated game systems, major computer companies, and audio chipset manufacturers licensing low-end virtual audio technology. As a result of increasing availability and the lower cost of the technology, virtual audio should become a common component of VE systems within the next five years.

5. PRIMARY USER INPUT INTERFACES

The majority of the types of interfaces discussed in this report are those that provide sensory inputs to the user, specifically, visual, auditory, and haptic inputs. It is important to look at the other side of the coin, that is, inputs from the user to a VE system. While trackers provide one type of input to a VE, this is usually a passive form of input as far as the user is concerned and not a means whereby the user can specify commands to the system. The devices discussed in this section provide the primary means for direct user interaction with VEs. They allow specifying input commands that serve, for example, to determine movement through the environment or effect virtual object manipulations. These devices are not limited to use with VEs, but are used with many computer-based applications that require interaction with 3-D objects, such as computer-aided design (CAD) applications.

The form of the input command itself varies from naturalistic hand gestures, menu item selection, or object selection, to the operation of buttons with preset functions. Two categories of devices are considered in this section: (1) gloves and exoskeleton devices that can support the first two type of command form, and (2) 3-D pointing devices that support menu item and object selection and often provide a number of user-programmable buttons. The choice of which type of device is preferable in any particular application largely rests on the degree of naturalness desired for the interface.

5.1 Whole-Hand and Body Inputs

Sturman and Zeltzer (1993) define whole-hand inputs as "the full and direct use of the hand's capabilities for the control of computer-mediated tasks... [providing an interface that] makes maximal use of the naturalness, dexterity, and adaptability of the human hand." Since one of the primary goals of VEs is to enable natural methods of interaction, it is hardly surprising that whole-hand input in the form of hand gestures are one of the most commonly used methods for providing user inputs in the VE. Gestures are used to transmit messages relating to desired movement through the environment, and to select and manipulate objects, even objects as diverse as option menus and soda cans. While the advantages of gestures can seem clear in many VE applications, there has been little experimental evaluation of the use of whole-hand input compared to other forms of input. In one set of experiments, Sturman and Zeltzer (1993) compared the use of gestures to conventional input via a set of dials for the control of a six-legged walking robot. The gesture input was found superior for control of low-level walking and for high-level manipulation, whereas the conventional input gave best performance when steering the robot: the gesture input gave the

best performance when the required interaction mapped well to the human hand in terms of naturalness, adaptability, and dexterity, and when task characteristics (such as required degrees of freedom, resolution, and steadiness) mapped well to hand action capabilities.

While the trackers previously discussed can provide information about absolute position of a user's hand in space and palm orientation, use of the hand as an input device generally requires additional information. Gesture recognition is dependent on information about relative finger positions and this is determined by measuring the angles of joints in the fingers. Similarly, many applications that require knowledge of the movement of various user body parts require information about the joints that control those body parts.

Before proceeding to look at specific products and current R&D in this area, this section presents applicable data on the range of motion of the human hand, arm, and shoulder, and human capabilities in sensing the positions of these joints.

5.1.1 The Human Hand and Arm Position Sense

As shown in Table 8, the range of motion provided by human joints varies quite widely. The human hand in particular is capable of great freedom of movement, providing 29 DOFs. Twenty-three of these DOFs are exhibited by the joints in the hand, and the remaining 6 DOFs by palm.

Table 8. Range of Motion for Hand, Arm, and Shoulder Joints^a

Joint	Motion	Range	Joint	Motion	Range
Thumb	Palmar Adduction	90°	Shoulder	Abduction/Adduction	150-184°
	Radial Abduction	80-90°		Media/Lateral Rotation	130°
	Opposition	90°		Horizontal Flexion/Extension	170°
	MCP Flexion	50°		Scapula Elevation/Depression	10-12 cm
	PIP Flexion	80°		Scapula Medial/Lateral Movement	15 cm
				Scapula Rotation	60°
Digits	Abduction/Adduction	±15°	Elbow	Elbow Flexion/Extension	145°
	Index MCP Flexion	86-90°		Forearm Supination/Pronation	155-180°
	Index MCP Extension	22-45°	Wrist	Flexion/Extension	85/170°
	Index PIP Flexion	100-110°		Radial/Ulnar Deviation	56-60°
	2nd finger MCP Flex.	91°		Abduction	15°
	2nd finger MCP Ext.	18°		Adduction	45-55°
	2nd finger PIP Flexion	105°			

a. Adapted from (Greene and Heckman, 1994), (Livingstone, 1983), (Reynier, 1993), and (Sturman, 1992).

With respect to position sensing, Tan et al (1994) report that the just-noticeable-difference (JND) for the proximal-interphalangeal (PIP) and metacarpal-phalangeal (MCP) finger joints is 2.5°, at the wrist and elbow 2°, and decreases to 0.8° at the more proximal shoulder joint.

5.1.2 Commercially Available Devices

There are four glove-based devices for whole-hand tracking on the market, and these are described below. Two exoskeleton devices for hand and arm measurement are also commercially available. At the present time, only one body suit instrumented to measure the angles of various limbs is commercially available. Table 9 summarizes some of the key features of these devices. A glove product that is available in Japan but not, as yet, in the U.S. is Nissho Electronics Corporation's SuperGlove. Another glove-based product is the PC PowerGlove, under development by the makers of the Mattel Power Glove, Abrams Gentile Entertainment, Inc., and intended to replace the original glove. The PC PowerGlove is scheduled to be released in the first quarter of 1996. It is designed for use with PC video games, and supports position and orientation tracking in 6 DOFs, finger position measurement, and tactile and sweat feedback. Additionally, some of the force feedback devices discussed in Section 6.2 have integrated joint position measurement capabilities.

Virtual Technologies plans to bring a body suit to market in the near future, and Paradigm Shift is developing both a glove and a body suit interface device. Greenleaf Medical Systems has acquired the licensing rights to market VPL's DataGlove and DataSuit for medical applications and are likely to acquire the remainder of VPL's assets. This would broaden their rights for use of the glove and suit technology in other application areas and may lead to the development of some commercial products.

5.1.2.1 5th Glove


Fifth Dimension Technologies' 5th Glove uses proprietary optical-fiber flexor technology sensors. Each finger of the glove is fitted with a sensor that measures the average flexure of the finger. In its latest release, the 5th Glove also includes a 2-axis tilt sensor that measures $\pm 60^\circ$ roll and pitch orientation of the user's hand and can be mounted in either the horizontal or vertical direction. This new sensor allows the glove to emulate a mouse or a baseless joystick. For its physical structure, the glove uses stretch lycra material with the flexor sensors mounted on the fabric and the tilt sensor attached by velcro. A small electronics box is fastened to the glove and mounted on top of the wrist. A specification for the device is given in Figure 52. The device interface is serial RS-232 (3 wire) at 19.2 kbaud (full duplex). The number of gloves that may be supported simultaneously is limited by the number of serial ports.

The interface package is supported by Windows and DOS-based software that enable installation, glove calibration, and graphical (using approximately 80 polygons) representation of a virtual hand. A gesture recognition program can be trained to identify certain hand positions using a least squares fit algorithm. A program called KineMusica that converts finger bend data to MIDI output and allows playing a variety of musical instruments is included with the package. Finally, a DOS application, with C++ source code, provides raw glove data and can be used to support development of device drivers for other applications. Additional software is provided with the latest release of the glove.

Table 9. Characteristics of Commercial Available Glove, Body Suit, and Hand/Arm Exoskeleton Input Devices

Device Name	Vendor	Device Type	Measurements	Sensor Resolution	Weight	Software Support	Price
5th Glove	Fifth Dimension Technologies	Glove	Finger flexure, hand roll & pitch	8 bit	~350 g	Windows & DOS interface package & drivers, C++ source code, gesture recognition, other	\$495
CyberGlove	Virtual Technologies, Inc.	Glove (18 sensor)	PIP, MCP finger joint angle, finger abduction, thumb opposition, palm arch, wrist flexion & abduction	0.5°	2.5 oz	VirtualHand Silicon Graphics interface, DOS demo interface. (GesturePlus package separate)	\$9,800
		Glove (22 sensor)	3 flex sensors and abduction sensor per finger, thumb opposition, palm arch, wrist flexion and abduction				\$14,500
		Exoskeleton	16 DOFs for fingers, 4 DOFs for thumb				Contact vendor
Pinch Glove	Fakespace, Inc.	Glove (2)	Identification of finger/thumb contact ^a	N/A	0.7 oz	Silicon Graphics & PC interface	\$2,000
Position Exoskeleton ArmMaster	EXOS, Inc.	Exoskeleton	Shoulder flexion/extension, shoulder ab/adduction, shoulder int/ext rotation, elbow flexion/extension, forearm supination/pronation	0.1°	<1 lb (on arm)	Silicon Graphics interface	Contact vendor
TCAS Glove	T.C.A.S. Effects Ltd.	Glove (8 sensor)	PIP finger joint angle, 2 thumb joint angles, palm movement	Unavailable	Unavailable	Support for display of data to screen or file, interfaces to some VE toolkits	Starts at \$7,000
		Glove (11 sensor)	As 8 sensor version, plus: MCP finger joint angle, wrist flexation, adduction, and abduction				
		Glove (16 sensor)	Angles for all finger and thumb joints, wrist movements				
TCAS Data Wear	T.C.A.S. Effects Ltd.	Body suit	Custom designed for measurement at 32 sites	Unavailable	Unavailable	Support for display of data to screen or file, interfaces to some VE toolkits	Starts at \$30,000

a. This system interprets contacts between any two or more fingers or thumbs (on both hands) as a particular gesture.



Specification	
Update Rate	Each sensor 200 Hz (max)
Sampling Rate	For full hand, 200 samples/sec (max)
Sensor Resolution	Flex and tilt sensors, 8 bit
Latency	5 ms
Calibration	Once per user
Glove Sizes	One size
Glove Weight	~350 g
Interface	RS-232

Photo courtesy of Fifth Dimension Technologies

Figure 52. 5th Glove

This new software includes both 16-bit and 32-bit DOS drivers, a diagnostic program, a Windows 32-bit DDL, upgradable to Windows '95. In addition, it includes a Microsoft mouse emulator that interprets the user tilting his hand as cursor movement commands. Drivers to support the use of the 5th Glove with the Sense8 WorldToolKit are available, and drivers for Vream, Inc.'s VRCreator and Avril are expected by the end of 1995.

Fifth Dimension Technologies is developing a standard for usage of the 5th Glove. This standard is intended to ensure that a glove user, working with different applications, will not have to learn different sets of gestures for each application. It defines, for example, the gesture required to perform body and hand rotation to the right in a VE. Once stabilized, this standard will be provided with the 5th Glove as the "Recommended Implementation," and the defined gestures will be supported as defaults for the gesture recognition program. Meanwhile, default hand gestures have been defined to emulate button clicks.

The current price for the 5th Glove system is \$495 for right-handed users, and \$535 for left-handed users. Volume discounts are available. In addition to being available from the developers, Fifth Dimension Technologies in South Africa and 5DT (Europe) in Surrey, UK, these products are available in the US from General Reality Company.

Fifth Dimension Technologies also offers an ultrasonic tracker (see Section 3.1.1.21) that can be used with the glove or as a head tracker.

5.1.2.2 CyberGlove

Virtual Technologies, Inc. markets the *CyberGlove*, an instrumented glove primarily designed for manipulation of 3-D objects in the company's *CyberCAD* virtual design environment. The glove is constructed of a 80/20 Nylon/Lycra blend for flexibility, with mesh in the palm area and underside of fingers for ventilation. (Fingertips are left open to allow ease of keyboard typing or other physical object manipulation.) The *CyberGlove* uses

up to 22 sensors to measure joint angles: 3 flex sensors and an abduction sensor per finger, and sensors to measure thumb opposition, palm arch, and wrist flexion and abduction. The glove has a software programmable switch and LED in the wristband that can be used to control program input/output capability; preprogrammed functions include a time-stamp and readout of glove status. The wristband also provides mounting provisions for either a Polhemus or Ascension 6 DOF tracking sensor. Further details are given in Figure 53.




Figure 53. 22-Sensor CyberGlove

The *CyberGlove* comes with *VirtualHand* software for a Silicon Graphics workstation which displays a 2500-polygon Gouraud-shaded graphic representation of the user's hand and finger movements, or a lower resolution 325-polygon hand model. An executable version of this software that can be used for calibration and demonstration purposes is available for PCs. The *CyberGlove* system is supported under the Silicon Graphics version of WorldToolKit, with support for other versions of this VE toolkit to follow in the near future. It is also supported by Division, Inc. in some of their toolkit products. For an 18-sensor glove, the *CyberGlove* with interface unit and *VirtualHand* executable software is priced at \$9,800; the custom 22-sensor *CyberGlove* is \$14,500.

An additional Virtual Technologies, Inc. software product to support use of the glove was released in Fall 1995. Called the *GesturePlus*, this package uses neural networks that users can train to recognize their own customized hand gestures (up to 255 different gestures are possible). These gestures are associated with user-selected symbols to allow their mapping to user-defined actions. The *GesturePlus* system comes with an interface unit that performs the necessary gesture recognition processing and provides a serial RS-232 interface for connection to a range of computer platforms. Its introductory price is \$3,500.

5.1.2.3 Dextrous HandMaster

The Dextrous HandMaster is available from EXOS, Inc. It is a exoskeleton device that uses Hall Effect sensors to provide measurement of the joint flexion for four fingers and thumb. It can be used for providing motion commands in a VE or teleoperation envi-

ronment, or for recording finger motion. An interface to AT-compatible machines is provided. Further details are given in Figure 54. This product is made to order and price information is not available.




Figure 54. Dextrous HandMaster

5.1.2.4 Pinch Glove

The most recently announced glove interface is the Pinch glove system introduced by Fakespace, Inc., based on a prototype developed by researchers at the University of Central Florida, Institute for Simulation and Training. Unlike the other glove interfaces discussed here, the Pinch glove system does not measure finger joint angles. Instead, gloves are worn on both hands and contact between any two or more fingers, or thumbs, completes a conductive path, allowing the definition of a variety of "pinch" gestures that an application developer can map actions against. Over 1,000 gestures are theoretically possible. The gloves are constructed of a stretchable fabric and contain an electrical sensor in each fingertip. Each glove has a back-of-hand mount to accommodate a spatial tracker. The user's point of interaction in the VE is represented by a 3-D cursor. Further details are given in Figure 55.

The interface system, called the Pinch Hand Gesture system, consists of gloves for the left and right hand, electronics interface, and controlling software for either PCs or Silicon Graphics workstations. The system is supported by Sense8's WorldToolKit and support for other VE toolkits is under development. The price of the system is \$2000. Additional, single gloves can be purchased at \$100 each.

5.1.2.5 Position Exoskeleton ArmMaster

The EXOS, Inc. Position Exoskeleton ArmMaster (EAM I) is a pre-cursor of the Force Exoskeleton ArmMaster (see Section 6.2.2.2). It provides 5 DOF passive sensing of the upper and lower arm through a shoulder mechanism that provides 3 DOF sensing and an elbow mechanism that provides 2 DOF. These mechanisms are modular and are avail-




Specification	
Sampling Rate	< 1 msec
Calibration	Not needed
Glove Sizes	1 size
Glove Weight	0.7 oz
Interface	RS-232

Photo courtesy of Fakespace, Inc.

Figure 55. Pinch Glove

able separately. The sensors employed in the device are precision conductive plastic potentiometers. Further details are given in Figure 56. This device is made to order and price information is not available.



Specification		
DOF	Shoulder	3
	Elbow	1
	Forearm Sup/Pro	1
Joint Motions	Shoulder flexion/extension	120°
	Shoulder ab/adduction	120°
	Shoulder int/ext rotation	90°
	Elbow flexion/extension	130°
	Forearm supination/pronation	150°
Sensor Resolution	0.1° (joint computation ~1.5°)	
Weight	<1 lb on arm	

Photo courtesy of EXOS, Inc.

Figure 56. Position Exoskeleton ArmMaster

5.1.2.6 TCAS DATAWEAR

T.C.A.S. Effects Ltd. market a range of body tracking systems based on their patented conductive elastomer sensor technology. These systems are based on a washable body suit that is available in five sizes and consists of a jacket (with or without gloves) and pants that can be attached by three positional zippers and stud fastenings. Both a Lycra suit intended for use as an undergarment and a neoprene modular over-suit are available. The glove itself has standard configurations of 8, 11, or 16 sensors. The eight sensor version measures joint angles for the PIP finger joints, the two thumb joints, and palm movement. The 11 sensor glove adds measurement of the MCP finger joints and wrist flexation, abduction, and adduction movements. Finally, the 16 sensor version measures angles for all finger

and thumb joints, and wrist movements. Additional sensors are attached to the body suit according to the joints to be monitored. Another product, a rigid face mask, provides for monitoring facial expressions, including lip movements, using 7 to 12 sensors.

The number of sensors that can be supported simultaneously is limited by the control unit, which currently provides up to 32 channels. An increase to 64 channels is expected in 1996. To reduce tethering requirements, each set of 8 channels is packaged into a single standard 25 core cable. The tethering limits the operational range of the bodysuit and glove to 10 m², although extension cables can be added as necessary.

The software support for TCAS DATAWEAR provides for screen display of collected data or saving the data to file. Interfaces to VE toolkits, such as WorldToolKit, and computer animation packages are being developed. The basic eight sensor glove is available for \$7,000. Body suits are custom-developed, with a 32-sensor suit starting around \$30,000.

5.1.3 Current Research and Development

By and large, there is little research on whole-hand and body tracking to discuss. Most of the ongoing research and development is regarded as highly proprietary and little information is publicly available.

One effort for which detailed information could not be acquired, though not necessarily for reasons of privacy, was John Fairley's development of a body suit. This work began at the University of Illinois at Urbana-Champaign, Advanced Digital Systems Laboratory, with the development of a bodysuit that uses fiber-optic cables to measure the bend or rotation of joints. The data is transmitted to a PC where user movements are modeled by a 3-D mannequin. In subsequent independent R&D, this bodysuit is being refined so that it requires less calibration, monitors additional joints, and is cheaper to produce.

5.1.3.1 Armstrong Laboratory

Researchers at the Virtual Environment Interface Laboratory (VEIL) at Armstrong Laboratory, Crew Systems Directorate, see VE technology as providing a flexible, multi-modal medium that can support a broad range of adaptive interface concepts. Their work has two broad goals. One is to develop multi-modal, adaptive user interface concepts within a cooperative agent framework. Here the researchers are concerned with the cognitive design of interfaces that can be used to provide novel interaction channels for VEs. The other goal is to establish a quantitative relation between human performance (sensory, perceptual, psychomotor, and cognitive) in VEs and hardware, software, and model world properties of VE systems. The objective is to produce design trade-off nomographs that quantify the relation between various aspects of user performance and VE system factors that design engineers can use in the development of specific applications.

To date, the majority of the work has focused on the second goal. Led by Dr. Robert Eggleston, the researchers at VEIL have already conducted several experiments. The objec-

tive of one of the first experiments was to determine whether a person using a virtual hand controller could achieve performance comparable to that obtainable with a physical control device. The virtual hand controller used was a standard issue Nomex flyers' glove with attached Ascension Bird tracker, and the physical control device was a spring-loaded, return-to-center isotonic joystick. A single-axis control task was used as the experimental task, specifically, the Critical-Instability Tracking Task. The VE was provided by a single-seat cockpit simulator. The major finding of the experiment was that comparable performance can be obtained from the two device types (Eggleston, 1993). Subsequent experiments have looked at the impact of VE system delays on tracking performance, the difference between device types in recovering from tracking errors, and impact of the biomechanics of virtual controllers versus those of physical devices in catering for a range of body sizes. Information on the results of these other experiments is not publicly available at the present time.

Currently, the researchers are engaged in a series of experiments that are designed to investigate different types of control movements using virtual and physical controllers. Other on-going work is concerned with identifying those particular characteristics of multi-modal VE interfaces that enable performance similar to that obtainable in the real world. Since earlier work has indicated that human perceptual systems are not engaged in a VE in the same way that they are in the real world, this area of work is likely to include basic investigations into perception issues.

Another area of current work is concerned with sensing whole-body movement. Here the researchers are exploiting the new techniques developed at MIT for electric field sensing. The general principle is that when capacitive coupling is established between a human actor and a generated electric field, because the coupling profile changes as a function of the location of the body parts, this profile can be used to track the movement of those parts. Researchers have developed a basic capability for this type of tracking and are currently engaged in further studies of the basic technology and its use.

Work supporting the development of design trade-off nomographs will continue. Starting in 1996, however, increased attention will be paid to multi-modal, adaptive user interface concepts.

5.1.3.2 Georgia Institute of Technology

In a study completed in 1994, researchers at Georgia Institute of Technology, Graphics, Visualization and Usability Center, performed an experimental evaluation of Virtual Technologies' CyberGlove, Model CG1801. The objective of this work was to determine the glove sensor characteristics and, on this basis, determine its suitability for person-independent gesture recognition. The three-part experiment investigated the level of sensitivity of the glove sensors, the performance in recognizing angles, and factors that affect accuracy of angle recognition. Sixteen subjects were selected who provided a range of hand sizes. As reported by Kessler, Hodges, and Walker (1994), the measurement range for each

joint varied considerably. For example, the average range for the subjects' MCP joints varied from 99.5° to 104.6° across the four fingers, and the average range for the PIP joints varied from 96.8° to 118.2° across fingers; this data was collected prior to any glove calibration. After simple calibration that measured two extreme values of flexion for each finger joint to allow for a translation from discrete values to joint angles, collected data showed that the angle calculated did represent the actual angle of the joint. Given a set of possible angles (that are related to particular gestures), the reported angle was correctly classified 90% of the time, with some notable exceptions (thumb joints, abduction joints, and high values of joint flexion). Additional calibration based on data already collected for an individual, or a group's data, served to increase angle recognition, although problems remained in recognizing finger abduction angles, and accuracy decreased as angles increased. Repeatability of the recognition was found to be dependent both on the joint and angle being measured. Device noise and hand size were not significant factors in recognition accuracy. Based on this evaluation, the researchers concluded that current technology supports only a limited gesture recognition ability.

Future work may include an investigation of the effect that a glove-based interface has on the sense of presence experienced by VE users.

5.2 3-D Pointing Input

3-D input via a mouse-like joystick or trackball device is a dominant form of user input in both immersive and non-immersive VEs. Such devices are used as a means of navigation, object selection, and, in some cases, object manipulation. They provide no feedback to the user, other than that which is inherent in the device's physical construction. They use acoustic, electromagnetic, inertial, mechanical, or optical transducers to convert a physical phenomenon, force or velocity, into a measurable signal.

5.2.1 Commercially Available Devices

There are fewer of these devices on the market than one might expect, and what products there are differ quite widely in physical form and functional capabilities. Six devices are reported below and their characteristics are summarized in Table 10. Information on additional products from Global Devices was not available. Two more devices that are essentially elements, or special cases, of more general tracking systems are discussed in Section 3.1; these are the Logitech 3D Mouse and the RPI Mouse-Sensor3D.

5.2.1.1 CyberWand

Specifically designed as a navigational device for VE applications, InWorld VR Inc.'s CyberWand is a handheld joystick. In its base format, the CyberWand provides 2 DOFs via a hat-shaped device that controls either left/right or forward/back movement, one of the four programmable buttons is used to specify which axis is being controlled. In general, the buttons can be used to achieve movement in 6 DOFs by specifying how movement of the hat sensor should be interpreted. Source code is provided for the CyberWand, allow-

Table 10. Characteristics of Commercially Available 3-D Point Input Devices

Product	Vendor	Input	Device Type	Software Support	Price
CyberWand	InWorld VR, Inc.	2-D ^a	Handheld joystick (4 buttons)	WorldToolKit	\$60-104
Immersion Probe-MD	Immersion Corporation	3-D	Desktop boom-based stylus	Interface drivers	\$1,995
Magellan 3D Controller	Logitech, Inc.	2- or 3-D	Desk-based mouse (9 buttons)	AutoCAD	\$550
RingMouse	Kantek, Inc.	3-D	Finger ring (2 buttons)	Windows 3.1 and 95, DOS interface	\$120
Spaceball 2003	Spacotec IMC Corporation	3-D	Desk-based mouse (8 buttons)	WorldToolKit, VREAM, Superscape	\$1,195
Space Controller	Spacotec IMC Corporation	3-D	Desk-based mouse (2 buttons)	WorldToolKit, AutoCAD, CADKey, Microstation	\$595

a. 3-D input available by attaching a Polhemus or Ascension tracking sensor.

ing users to develop their own device driver, although special drivers for Sense8's WorldToolKit are included. In addition, a demonstration program is provided, example programs for WorldToolKit, and a calibration program for a ThrustMaster dual-port ACM card. With the ThrustMaster ACM card, the CyberWand is available for \$99, by itself it costs \$60. Specification information for this device is not available.

An alternate version of the CyberWand provides special provisions for the attachment of a Polhemus or Ascension tracking sensor. The attachment of such a sensor provides true 6 DOF control for the CyberWand, leaving the four buttons open for other uses. In this case, the hat-sensor can be used to achieve large movements that are beyond the reach of the handheld CyberWand. This version of the CyberWand costs \$104.

5.2.1.2 Immersion PROBE-MD

Immersion Corporation has recently released the Immersion PROBE-MD, a 3-D input device intended for use with VE, CAD, telerobotic, medical imaging, and graphics applications. The PROBE-MD is a mechanical arm controller with a stylus tip that can be freely moved in 6 DOFs, further details are provided in Figure 57. The price of PROBE-MD is \$1,995. A Developer's Programming Library (in the C programming language) is available for PC, Mac, SGI, or Unix platforms for an additional \$175.

The probe's interface is capable of supporting an additional input device. A digital foot pedal (standard or heavy duty), digital hand switch, or analog foot pedal are available for use with this spare channel.




Figure 57. Immersion PROBE-MD

5.2.1.3 Magellan 3D Controller and Space Controller

Logitech, Inc. is phasing out their CyberMan joystick, and has introduced the Magellan 3D Controller to replace it. An early version of this controller was used onboard the space shuttle Columbia for space-based telerobotics applications.

The Magellan 3D Controller provides a spring-mounted puck for controlling movement and absolute measurement is achieved by a patented linear optical measuring system. The device can be used in either 2-D or 3-D mode. For 3-D application, the controller is operated by the user to position an object, while working on that object with an ordinary mouse. The controller provides nine user-programmable buttons, some of which can be used to adjust sensitivity and motion control (for example, only reporting the coordinate with the greatest magnitude, or only rotation coordinates). Device drivers are available for PC, IBM, Sun, HP and Silicon Graphics platforms. The PC version is additionally supported with an AutoCAD driver and demonstration software. Further details are given in Figure 58. The Magellan 3D Controller is available for \$550.

5.2.1.4 RingMouse

Kantek, Inc. recently announced a new type of 3-D mouse, called the RingMouse. As its name suggests, this device takes the form of a ring worn on the user's finger. The ring itself consists of an ultrasonic transmitter, held in position by a velcro strap. This transmitter sends both an infrared and ultrasonic pulse to a receiver that is mounted on top of a monitor, the delay between receiving these signals is used in determining the ring position. The system has a 3 ft tracking area and is wireless. The ring is powered by a long-life watch battery and an automatic sleep mode switches the power off if the ring is unused for a minute or so, until one of the buttons is pressed. The ring itself has two programmable buttons, intended to be pressed using the thumb. Since the RingMouse receiver plugs into the




Figure 58. Magellan 3D Controller

serial port, this device can be used in conjunction with a traditional mouse. Further details for the RingMouse are given in Figure 59.





Figure 59. RingMouse

The software support for RingMouse includes an emulator for a traditional mouse. There is also a joystick emulator for use with DOOM and other DOS-based games, and special interfaces for the Dark Forces and Descent games. A 3-D hoop toss game is provided with the RingMouse. The price for RingMouse is approximately \$120.

5.2.1.5 Spaceball 2003 and Space Controller

Spacotec IMC Corporation developed and market the Spaceball 2003 controller that is designed to support interactive control of 3-D models and VE navigation in simultaneous 6 DOFs. A patented sensing technology is embodied in a PowerSensor ball that can be pulled, pushed, or twisted using the fingertips to control movement. Eight buttons provide motion control filters, performing functions such as switching rotations on/off, adjusting sensitivity, and view reset. Further details are given in Figure 60. The price of Spaceball 2003 is \$1,195.

The proprietary SpaceWare IMC interface software supports interfaces with many CAD and computer-aided manufacturing applications. Spaceball 2003 is also supported in a number of VE world building packages, including WorldToolKit, VREAM, and Super-scape. A Software Developer's Kit is available to provide the source code and information needed to integrate Spaceball support into custom applications.



Specification	
Max/Min Detectable Force	~4.5/0.1 lb
Spatial Resolution	10 bits
Max/Min. Detectable Torque	~5.3/0.1 lb/in
Dimensions	9.1 x 3.5 x 4.9 in
Weight	1.75 lb
Interface	RS-232, 9 pin

Photo courtesy of Spacotec IMC Corporation

Figure 60. Spaceball 2003

Specification	
Max/Min Detectable Force	0.1/4.6 lb
Max/Min. Detectable Torque	0.1/5.3 lb/in
Resolution	10 bits
Dimensions	7.6 x 3 x 4.4 in
Weight	1.25 lb
Interface	RS-232, 9 pin

Figure 61. SpaceController

A more recent product, SpaceController, provides a lower cost version of the Spaceball 2003. Here only two buttons are provided and these control pop-up menus that give access to the functionality of the Spaceball 2003 buttons. Further details for the SpaceController are given in Figure 61. Its price is \$595.

An additional Spacotec IMC Corporation 3-D input product, the Spaceball Avenger, is intended only for use in PC video games and not discussed further.

5.2.2 Current R&D

Here, again, there are few research and development efforts to report. In part, this can be attributed to the fact that there are no major outstanding technical issues in the construction and operation of this type of 3-D input devices. However, data on the utility of the particular devices for specific types of manipulations is still needed. More generally, as the example of the work underway at the University of Toronto shows, there are many human factors and cognitive engineering concerns that remain to be addressed.

5.2.2.1 Digital Image Design Inc.

Digital Image Design Inc. is designing a 3-D input device that will be the successor to their Cricket device, shown in Figure 62, that is no longer under manufacture.



Photo courtesy of Digital Image Design, Inc.

Figure 62. Cricket

While still in early design stages, the new device is expected to be functionally similar to the Cricket. That is, it will be a device specifically designed to support VE applications that require substantial free-space manipulation, probably providing buttons to support an occasional need for 2-D operations (such as menu manipulation), object picking, and object grabbing. Current work is focusing on ergonomic issues, particularly in the areas of reducing the stress placed on a user's hand. No dates for the expected release of the new product are yet available.

5.2.2.2 University of Toronto

For the past several years, researchers at the University of Toronto, Department of Industrial Engineering, have been investigating human factors issues concerned with 6 DOF input techniques for the manipulation of objects in 3-D environments. This work is being led by Dr. Shumin Zhai and Dr. Paul Milgram. The overall goal of the work is to determine critical factors for the design of 6 DOF input devices and their impact on human manipulation performance.

A central aspect of the work has been the development of a model for classifying 6 DOF input devices along the human factors dimensions of mapping (position versus rate control), sensing mode (ranging from isotonic, through elastic, to isometric), and degree of integration (based on number of discrete controls to be manipulated). This model has served as the framework for a series of experimental studies. In the first experiment, the researchers compared isotonic-position, isotonic-rate, isometric-rate, and isometric-position control approaches (Zhai, 1993). A glove was used for the first two approaches and a spaceball for the third and fourth. In an experimental 6 DOF docking task presented via a non-immersive VE system, using eight subjects, the researchers found a strong interaction between sensing mode and mapping, and that isotonic-position and isometric-rate approaches gave the best performance.

While the first experiment showed that the isotonic-position device was more direct, it was tiring to use. The isometric rate device was less fatiguing but provided little kinesthetic feedback to the user. Consequently, the objective of the second experiment was to assess an intermediate approach by comparing elastic-rate and isometric-rate control approaches. (The elastic-rate device developed by the researchers is called the EGG, further details about this device are given in Figure 63 below.) The same experimental task was used, this time with twenty-six subjects. Significant differences in performance were found only during early stages of learning with the devices, with the EGG outperforming the spaceball. The third experiment was designed to see whether larger differences would result

with a more challenging experimental task. Again using the EGG and spaceball, twenty six subjects were asked to perform a pursuit tracking task. As in the previous experiment, subjects' performance with each device improved substantially with practice. Zhai and Milgram (1993) report that the elastic-rate EGG produced better scores, especially in the early learning stage.




Figure 63. EGG

The data collected in the third experiment were also analyzed to determine the subjects' spatial accuracy in the x , y , and z directions. The mean tracking error in each direction was found to decrease over time, but the mean error in the z direction was significantly greater than in the x direction through all experimental phases. The mean error in the y direction was comparable to that in the z for early training and subsequently greatly decreased to the level of the mean error in the x direction. Overall the mean error in the z direction was 20% and 40% greater than that in the y and x directions, respectively. These findings were the same for both types of devices. The researchers hypothesize that a reason for the superior performance in tracking in the x direction than the y direction might be a higher attentional resource priority for horizontal movement.

A fourth experiment studied the issue of which joints and muscle groups should be used for 6 DOF manipulation. For this experiment, two isotonic-position control techniques were tested in a 6 DOF docking task. One technique utilized the user's wrist, elbow, and shoulder, while the other technique additionally made use of the user's fingers. The results showed that the participation of fingers significantly improved the task performance.

In a final experiment to be reported, issues concerning the visual representation formats of users' input control actions in relation to a target object were investigated. In a 3-D dynamic target acquisition task, it was found that both binocular displays and partial occlusion through semi-transparency, a novel graphic technique, were beneficial. In particular, the use of semi-transparent surfaces appeared to enhance human performance in discrete tasks more than the classical stereoscopic viewing technique.

Currently, these researchers are investigating human behavior in coordinating hand movements in 3-D environments. This work includes developing new measures for correlations and determining how human coordination relates to different interface designs.

5.3 Summary and Expectations

In the case of glove-based devices, this is an area of current growth. All the glove-based devices on the market are relatively new products and several additional products are expected to be released in the near future. The motivation behind these devices is to allow users to manipulate a VE in a similar manner to that in which they would manipulate a real environment. The current set of commercially available products do provide this capability but in a limited manner. The limitations arise chiefly from the lack of sensory feedback to the user's hand and the inability for fine discrimination between gestures. With respect to this last issue, technical improvements can be expected to occur in the next few years, but it seems likely that significant improvements in gesture recognition are more likely to result from a context-based approach for gesture recognition, a topic that does not seem to be receiving attention. Even so, the common use of gloves as a primary VE interface device is expected to continue.

The alternative to glove-based devices, exoskeleton devices, are expensive and encumbering. While some products are on the market, these are built to order and intended to be tailored to particular applications where precise joint measurement is required. In the case of VEs, these applications will be the exception rather than the rule and it is highly unlikely that exoskeleton devices will come into widespread use. Other uses of these devices include various specialized medical applications, these also are unlikely to provide a large market demand.

There are no particular technical challenges in the design and manufacture of 3-D input devices and there are several general-purpose devices available as commercial products. Are more devices needed, that is, do current products provide the necessary functionality and quality? There is no evidence that, in general, user needs are not being met. Consequently, while new devices may become available in the next few years, either as totally new offerings or replacements for existing products, there is no reason to expect that the overall situation will change and large numbers of 3-D input devices will appear.

An area that is receiving some, but not enough attention, is consideration of the human factors issues in the use of these different types of devices. This concerns more than ergonomic design issues. Basic questions pertaining to the usability and appropriateness of different device types need to be answered. Such questions should include consideration of both the application and the characteristics of the other types of display that are being used, primarily visual and auditory displays. They should consider not only the requirements a device places on motor skills, but also any cognitive burden that takes human processing resources away from the primary application task. While there is not a lot of evidence to suggest that these issues will receive in-depth investigation, hopefully the next few years

will see more work in this area. Meanwhile, there is a lack of data on the comparative capabilities and usability of current devices that can help users in selecting one over another.

6. HAPTIC INTERFACES

The human haptic system has an important role to play in human interaction with VEs. Unlike the visual and auditory systems, the haptic sense is capable of both sensing and acting on the environment and is an indispensable part of many human activities. In order to provide the realism needed for effective and compelling applications, VEs need to provide inputs to, and mirror the outputs of, the haptic system. Inputs to the haptic system are in the form of haptic displays and outputs are motor action commands, where the primary input/output variables are displacements and forces.

Haptic sensory information is distinguished as either tactile or kinesthetic (sometimes called proprioceptive) information. The difference between these is best illustrated by example. Suppose the hand comes up to an object suspended in space. The initial sense of contact is provided by the touch receptors in the skin, which also provide information on the contact surface geometry, the surface texture of the object, and slippage. When the hand applies more force, kinesthetic information comes into play providing details about the position and motion of the hand and arm, and the forces acting on these, to give a sense of total contact forces, surface compliance, and (if the hand is supporting the object in some way) weight. In general, of course, tactile and kinesthetic sensing occur simultaneously.

In order for the hand to manipulate the object, say move it horizontally, rotate it, or pinch it, the haptic system must issue motor action commands that exert forces on the object. These forces are highly dependent on the type of grasping that is used. Power grasping employs all the fingers and the palm, whereas precision grasping uses only the fingertips. Which is appropriate in a specific circumstance depends on such factors as the forces to be exerted and the dexterity of manipulation required. The manner in which the object being manipulated responds depends on the laws of physics and, potentially, a host of other sciences. That response, however, will be signalled by the haptic senses and may, in turn, guide further manipulation.

These, then, are the types of capabilities desirable for VEs and topics of this section. However, the following discussions are limited to consideration of the human hand, arm, and torso. Issues pertaining to whole body movement are addressed in Section 7.

6.1 Tactile Interfaces

As indicated above, tactile sensing plays an important role in object discrimination and manipulation. In many situations, it is either indispensable or critical for task performance. For example, there are reports of surgeons performing laparoscopic surgery who

find the visual feedback of laparoscopic instruments insufficient and insert a finger into the skin opening to feel for the presence of tumors in underlying body tissue. The fact that the lack of tactile feedback makes certain tasks more difficult is substantiated by experiments performed by Massimino and Sheridan (1993). Additional experiments have demonstrated the value of tactile feedback, for simple tracking tasks (Patrick, Sheridan, and Massimino, 1990), for reaction time reduction in target pointing (Akamatsu, 1994), and in degraded visual conditions (Massimino and Sheridan, 1993). While the presence or absence of tactile sensing undoubtedly has an impact on the sense of immersion experienced by VE users, there are no known studies that have investigated this.

In addition to the obvious example of sensing environment or object temperature, tactile sensing can support many discrimination activities that force sensing cannot. Tactile sensing, for example, is needed to determine the local shape and texture of objects and for detecting slip. It also provides information on surface compliance, elasticity, viscosity, and electrical conductivity. While the tactile ability to sense vibration is critical for determining surface texture, it is also valuable in its own right. Sensing of high frequency vibrations is a major component of many tasks and, in some cases, detection of vibration can be the goal of the work. Kontarinis and Howe (1995), for example, have shown that the presentation of high frequency vibrations can enhance performance of certain tasks by reducing reaction times or permitting minimization of applied forces. Since force feedback does not occur prior to any surface deformation, tactile sensing is also required for initial contact detection. Massimino and Sheridan (1993) have shown that tactile feedback can provide a significant performance improvement over force feedback for detecting the presence of contact forces, and tactile feedback provides similar, or superior, performance for detecting the magnitude of contact forces and for tracking a sustained contact force. Lastly, Howe (1992) has shown that tactile feedback is a necessary support to force feedback when gauging the minimum forces necessary for precise manipulation tasks.

It is also useful to note that, in some circumstances, one type of tactile display can be substituted for another. For example, Ino et al (1993) showed that temperature displays can be used to support discrimination of object materials and Morgan (1965) uses such a display to create the sensation of pressure or object contact.

At the current time, no known VEs in practical use support a tactile interface (the term "practical use" refers to systems either commercially available or those in everyday use by users, as opposed to developers.) By default, these systems use visual and/or auditory senses to substitute for the tactile sense; for example, by sounding an auditory tone when the user comes "in contact" with a virtual object.

Tactile stimulation can be achieved in a number of different ways. Those being used for VE systems include mechanical pins activated by solenoid, piezoelectric crystal, and shape-memory alloy technologies, vibrations from voice coils, pressure from pneumatic systems, and heat pump systems. The major strengths and weaknesses of these different approaches are summarized in Table 11. Other technologies, such as electrorheological

Table 11. Tactile Feedback Actuator Technologies^a

Technology	Description	Advantages	Disadvantages
Piezoelectric crystals	Changing electric fields causes expansion and contraction of crystals	- High spatial resolution	- Restricted to resonant frequency
Pneumatic	Takes many forms. As air-jets, provides an array of air nozzles that can be gated to a display pattern. As air-rings (cuffs), like miniature blood pressure cuffs. As bladders (bellows), often the size of a finger pad and held against the finger by a glove or band. As an array of tiny pressurized bladders, many to a single finger pad.	- Low mass on hand	- Poor spatial and temporal resolution - Limited bandwidth
Shape Memory Alloy	SMA wires and springs contract when heated and expand again as they cool under stress	- Good power-to-mass ratio	- Low efficiency during contraction - Heat dissipation problems limit relaxation rate of wires
Solenoid	Magnetic coil applies force to ferrous plunger	- High steady-state forces - Better bandwidth than other materials (except for piezoelectric crystals and voice coils)	- Relatively heavy - Nonlinear, can require extra effort to control
Voice coil	Voice coil vibrates to transmit low amplitude, high frequency vibrations to the skin.	- High temporal resolution - Relatively small, does not obstruct normal movement ranges of the fingers	- Poor spatial resolution - Limited scalability
Heat pump	Solid state device that moves thermal energy to heat or cool the skin	- No fluids required	- Poor spatial and temporal resolution - Bulky - Limited bandwidth

a. Based on Hasser (1995, 1996)

fluids that harden under the application of an electric field also are under investigation. Additional technologies found in medical applications, such as electrotactile and neuromuscular stimulation, have not yet been used and their invasive nature makes future use unlikely.

6.1.1 The Human Tactile Sense

There are four kinds of sensory organs in the hairless skin of the human hand that mediate the sense of touch. These are the Meissner's Corpuscles, Pacinian Corpuscles, Merkel's Disks, and Ruffini Endings. As shown in Table 12, the rate of adaptation of these receptors to a stimulus, location within the skin, mean receptive areas, spatial resolution, response frequency rate, and the frequency for maximum sensitivity are, at least partially, understood. The delay time of these receptors ranges from about 50 to 500 msec.

Table 12. Functional Features of Cutaneous Mechanoreceptors^a

Feature	Meissner Corpuscles	Pacinian Corpuscles	Merkel's Disks	Ruffini Endings
Rate of adaptation	Rapid	Rapid	Slow	Slow
Location	Superficial dermis	Dermis and subcutaneous	Basal epidermis	Dermis and subcutaneous
Mean receptive area	13 mm ²	101 mm ²	11 mm ²	59 mm ²
Spatial resolution	Poor	Very poor	Good	Fair
Sensory units	43%	13%	25%	19%
Response frequency range	10 - 200 Hz	70 - 1,000 Hz	0.4 - 100 Hz	0.4 - 100 Hz
Min. threshold frequency	40 Hz	200-250 Hz	50 Hz	50 Hz
Sensitive to temperature	No	Yes	Yes	At > 100 Hz
Spatial summation	Yes	No	No	Unknown
Temporal summation	Yes	No	No	Yes
Physical parameter sensed	Skin curvature, velocity, local shape, flutter, slip	Vibration, slip, acceleration	Skin curvature, local shape, pressure	Skin stretch, local force

a. Adapted from (Shimoga, 1993b), (Bolanowski et al, 1988), (Kontarinis, 1993), and (Reynier and Hayward, 1993).

It is important to note that the thresholds of different receptors overlap, and it is believed that the perceptual qualities of touch are determined by the combined inputs from different types of receptors. The receptors work in conjunction to create an operating range for the perception of vibration that extends from at least 0.04 to greater than 500 Hz (Bolanowski et al, 1988). In general, the thresholds for tactile sensations are lowered with increases in duration. Skin surface temperature can also affect the sensitivity of sensing tactile sensations.

These details provide some initial guidance for the design and evaluation of tactile display devices in such areas as stimulus size and duration, and signal frequency; perhaps

constraining the type of display technology used. For example, Kontarinis and Howe (1995) note that the receptive areas and frequency response rates indicate that a single vibratory stimulus for a fingertip can be used to present vibration information for frequencies above 70 Hz, whereas an array-type display might be needed for the presentation of lower frequency vibrations.

Additional information is available when looking at a higher level that the receptors just discussed, that is, at the receptivity of the skin itself. The spatial resolution of the fingerpad is about 0.15 mm, whereas the two-point limen is about 1 to 3 mm. Detection thresholds for features on a smooth glass plate have been cited as 2 μm high for a single dot, 0.06 μm high for a grating, and 0.85 μm for straight lines. Researchers have also looked at the ability to detect orientation. The threshold for detecting the direction of a straight line has been measured at 16.8 mm. When orientation is based on the position of two separate dots, the threshold was 8.7 mm when the dots were presented sequentially, and 13.1 mm when presented simultaneously. Reynier and Hayward (1993) discuss these findings and the results of additional work in this area. Data on the temporal acuity of the tactile sense is also reported by these researchers, who note that two tactile stimuli (of 1 msec) must be separated by at least 5.5 msec in order to be perceived as separate. Although, in general, increases in tactile stimulus duration can lower detection thresholds.

In a set of psychophysical experiments that investigated the capability of the human fingertip to detect strain, Ino (1993) found that the stimulus threshold was highly dependent on the motion of the skin contact surface (velocity, direction, and rotation), and surface viscosity and temperature, though not greatly affected by surface roughness. The reported findings are shown Table 13.

Table 13. Stimulus Thresholds for Strain

Stimulus	Threshold	Rate of Change
Velocity	50 μm at 0.2 mm/sec decreasing to 20 μm at 4 mm/sec. Threshold is related to direction as longitudinal < slant < transversal (in particular, threshold at 1 mm/s in longitudinal direction is ~ 0.5 that in trans. direction)	-13 dB/dec transversal, -11 dB/dec slant, -6 dB/dec longitudinal
Rotation	0.046° at angular velocity of 21.6°/sec, increasing to 0.265° at 0.72°/sec	-9 dB/dec
Surface Viscosity	500 μm displacement detectable at 233 $\mu\text{m/s}$ for viscosity of 500 cSt, decreasing to detection at 986 $\mu\text{m/s}$ for 30,000 cSt	6.9 dB/dec
Surface Roughness	35 μm for average grain size of 75 μm , increasing to 41 μm as grain size decreases to 8.5 μm	-
Surface Temperature	Minimum when contact surface near 32°, increasing with temperature	-

Burdea and Coiffet (1994) have summarized what is known about the human hand sensing bandwidth, as reproduced in Figure 64. Additional details are available with respect to vibration. The human threshold for detection of vibration at about 28 dB (relative to 1 μm peak) for frequencies in the range 0.4 - 3 Hz, this decreases for frequencies in the range

of 3 to about 250 Hz (at the rate of -5 dB/octave for the range 3 - 30 Hz, and at a rate of -12 dB/octave for the range 30 - 250 Hz), for higher frequencies the threshold then increases (Shimoga, 1993b).




Figure 64. Human Hand Sensing Bandwidth

Reprinted by permission of John Wiley & Sons, Inc. and IEEE.

The literature also provides information on the just-noticeable-difference (JND) for changes of temperatures. Researchers Yarnitsky and Ochoa (1991) conducted experiments that looked at the JND of temperature change on the palm at the base of the thumb. They found that two different measurement methods gave different results, and the difference between results increased as the rate of temperature change increased. Using the more traditional measurement approach based on a method of levels, and starting at a baseline temperature of 32°C, the rate of temperature change (1.5, 4.2, and 6.7°C/sec) had no detectable

effect on the JND for warming temperatures ($\sim 0.47^\circ$) or cooling temperatures ($\sim 0.2^\circ$). Subject reaction time was independent of the method used, and also independent of the rate of temperature change, although the reaction time for increases in warming ($\sim 0.7^\circ$) was significantly longer than the reaction time for increases in cooling ($\sim 0.5^\circ$). In reviewing work in this area, Zerkus et al (1995) report on findings that the average human can feel a temperature change as little as 0.1°C over most of the body, though at the fingertip a sensitivity of 1°C is typical. He also states that the human comfort zone lies in the region of 13 to 46°C . LaMotte (1978) reports that the threshold of pain varies from 36 to 47°C depending on the locus on the body, stimulus duration, and base temperature.

6.1.2 Commercially Available Interface Devices

Currently, few tactile interface devices are commercially available. EXOS, Inc. market the Touchmaster that is intended to present a sense of object contact to the user. It can be used independently or with their Dextrous HandMaster (see Section 5.1.2.3 for a discussion about the Dextrous HandMaster). Additionally, EXOS, Inc. sells a Hand Exoskeleton Haptic Display (HEHD) that provides both tactile and force feedback displays to the hand. Since the HEHD uses the same tactile display as the Touchmaster, it is not discussed in detail here, but more information is available in the section on commercially available force feedback interface devices (Section 6.2.2.6). Xtensory, Inc. market the Tactool system that also provides object contact feedback. A very different device, the Displaced Temperature Sensing System (DTSS), is marketed by CM Research, Inc. to provide temperature feedback. The characteristics of these products are summarized in Table 14 and they are discussed in more detail in the following subsections.

Until recently, Intelligent Systems Solutions, UK, (formerly the Advanced Robotics Center), marketed a tactile interface device based on pneumatic technology. A multi-channel pneumatic controller that included a pump, reservoir, and proportional pressure control channels was used to inflate air pockets that were designed to mount on existing, commercially available gloves. The resulting product was called the Teletact Glove and provided both tactile and force sensing. In its final version, the Teletact Glove used 30 air pockets with 2 pressure ranges. Twenty nine of these air pockets were positioned along the fingers and capable of a maximum pressure of 15 psi. The remaining air pocket was positioned in the palm and capable of 30 psi. Problems such as the deterioration of the air pocket material over relatively short periods, and a change in company focus, led to the product being withdrawn from the market.

Finally, for those interested in developing their own tactile interface device, or just experimenting with the tactile technology, Xtensory, Inc. and TiNi Alloy Co. both market tactile display kits based on SMA technology. Xtensory's Tactool Experimenter product, priced at \$250, consists of a single tactor with parts and a circuit diagram, and assembly is required. TiNi Alloy's Tactor Demonstration Kit provides a single $9 \times 20 \times 2.5$ mm tactor, pocket-sized Driver Box, interface cable for a PC-type serial port, and demonstration program. This tactor uses Muscle Wires constructed of Nitinol. With less than 1 volt, the

Table 14. Characteristics of Commercially Available Tactile Displays

Product	Vendor	Position of Display(s)	Type of Actuator	Effect	Tactile Sensation	Price
CyberTouch	Virtual Technologies, Inc.	4 fingers, thumb, and palm	(Proprietary)	Vibration (0-200 Hz)	Object contact	\$14,800
TouchMaster	EXOS, Inc.	4 fingers and thumb	Voice coil	Vibration (210-240 Hz)	Object contact	Contract vendor
Tactool System	Xtensory, Inc.	2 fingers	Blunt pins driven by SMA	Impulsive (30 g), vibration (20 Hz)	Object contact	\$1,500
Displaced Temperature Sensing System	CM Research, Inc.	Via thimble	Thermoelectric heat pump	Temperature change	Heating/cooling	\$10,000

tactor can pulse up to 2 cycles/sec on the skin to indicate contact with a virtual object, and variable pulse rates can be used to provide a sense of force feedback. The Tactor Demonstration Kit is priced at \$178.

6.1.2.1 *CyberTouch*

Released only in December 1995, the *CyberTouch* product from Virtual Technologies, Inc. provides a tactile feedback option for the *CyberGlove* (see Section 5.1.2.2). Tactile stimulators are attached to each fingertip and the user's palm to provide pulses or sustained vibration; they can be used individually or in combination to produce synchronized tactile patterns. The frequency of vibration generated is under user control and ranges from 0-125 Hz. Virtual Technologies, Inc. has applied for a patent on the actuator technology used and details of how the feedback is generated are not presently available. A photograph of Virtual Technologies' *CyberTouch* vibrotactile feedback option for the 18-sensor *CyberGlove* is shown in Figure 65, along with some specification details.




Figure 65. *CyberTouch*

A separate product, the *VirtualHand Toolkit*, provides a library of software routines that support use of *CyberTouch* including, for example, routines that update the stimulator actuators. *CyberTouch* itself comes with seven demonstration programs (including source code) that provide different force patterns. An additional demonstration shows the use of the tactile feedback in manipulating two balls, one suspended from a pendulum and the other resting on a simulated beach. *CyberTouch*, with a *CyberGlove*, is available for an introductory price of \$14,800. Upgrade or trade-in options are available for users who have previously purchased a *CyberGlove*. Also as part of the introductory promotion, the "Glove Mate" program provides a 34% discount towards the purchase of a second *CyberTouch* glove to allow tactile feedback to both hands.

6.1.2.2 *TouchMaster*

The EXOS, Inc. *TouchMaster* provides a tactile display to the tips of all four fingers and thumb using voice coil actuators. These actuators provide vibrotactile feedback that can be used to represent information about object contact. The voice coils are mounted on

a cable assembly and attached to the fingertips using velco bands, and driven by a signal condition box that can be interfaced to PC, VME, or other standard digital I/O busses. The standard configuration provides a fixed frequency of about 210 - 240 Hz at a constant amplitude, but optional variable frequency and amplitude electronics are available. Further details are given in Figure 66. The Touchmaster is built to order and price information is not available.




Figure 66. TouchMaster

In one experiment that investigated the effectiveness of the tactile display, researchers at EXOS, Inc. compared the use of a visual display (presented on a PC screen), the TouchMaster tactile display, and both visual and tactile displays in a task where subjects were asked to minimize the error in positioning of both thumb and index finger relative to the opposite sides of a virtual wall. The thickness of the wall was varied at 1.0 or 0.5 Hz between 30 mm and 80 mm. Data was collected from about 200 trials with each display combination and analyzed using a pairwise t-test. The results showed that trials with the visual display alone achieved an average error about two-thirds that produced with the tactile display alone, but the combination of both visual and tactile displays gave a performance increase over the visual display alone. Tactile feedback provided a five-fold improvement in the mean tracking error compared to the performance without tactile feedback.

In a second experiment, the difference between tactile and force displays as adjuncts to a visual display were investigated. (A 6 DOF Argonne E-2 master-slave manipulator was used to provide force feedback.) In this experiment, the subjects were asked to tap two targets alternately, as quickly as possible. The targets were of fixed width and set apart at three different distances, and an index of difficulty for each task was calculated based on this distance. For each index of difficulty, the results showed that the addition of either tactile or force displays to the visual display reduced task time by about one third. (The benefit of force feedback over that of tactile feedback depended on the task difficulty, ranging from 2% for the easiest task to 19% for the most difficult task.)

6.1.2.3 Tactool System

Xtensory, Inc. market the Tactool System where, again, the tactile display (called a tactor) is attached to a cable assembly for mounting on fingertips. While customizations are available, the base product consists of a single tactor (with associated cables) and a controller. The controller, which can support up to 10 tactors, provides the interface allowing the application software to give commands to "fire" a tactor. The primary interface is serial EIA232, but parallel, analog and MIDI interfaces are available to support such functions as reading sensors or daisy-chaining multiple Tactool Systems. Details on the tactor, Tactor Model XTT1, are given in Figure 67. This figure also shows a photograph of the Tactool System, together with a single pad-mounted tactor. The base Tactool System is priced at \$1,500, additional tactors are available at \$100 each. In addition, input sensors, which can be switches or force sensors, can be used for telerobotic applications.




Figure 67. Tactools System

Xtensory also market a 5 x 6 pin tactor array not intended for fingertip use.

6.1.2.4 Displaced Temperature Sensing System

The only known commercial product that provides temperature feedback for VEs is the Displaced Temperature Sensing System (DTSS) marketed by CM Research, Inc. Specifically designed for VE applications, when used with some tracking device, this system allows a temperature appropriate for the user's location in the VE to be sensed by the user's fingers. Using a thermode (an assembly of a thermoelectric heat pump, temperature sensor, and a heat sink), DTSS takes feedback from the sensor and regulates the temperature of the thermode surfaces.

The current product, DTSS Model X/10, is intended as a research tool. It consists of a controller, eight thermodes and connecting cabling. The controller can support eight thermode channels, each of which can be programmed as an input or output channel. It can be operated directly from the controller unit or, via a serial interface, by computer. Analog

inputs can be accommodated to allow tracking signals from external devices. The control law used for closed loop control of thermode temperature is the Proportional Integral Derivative law and the gains of each component in the law are adjustable. Safety features include large surface area heat sinks (the temperature of a heat sink is prevented from exceeding 40°C), and both a non-computer safety circuit and redundant software. DTSS is available for \$10,000, with additional thermodes priced at \$600 each. Further details are given in Figure 68. The initial version of DTSS, as shown in the figure, used velcro bands to attach ther-




Figure 68. Displaced Temperature Sensing System

modes to the user's fingers. This has been replaced in the current system by a thimble-type unit that the user can insert his finger into. The system comes with demonstration software (including source code) for PC and Mac platforms.

CM Research is currently developing a further version of DTSS that will use liquid cooled thermodes. These thermodes will provide better heat dissipation and, hence, support more rapid changes in temperature feedback.

6.1.3 Current Research and Development

Until the last few years, the majority of research and development on tactile interfaces focused on the development of reading aids for the visually impaired (see, for example, (Sherrick, 1984), (Shimizu, 1986), (Barfield and Furness, 1995)), tools to support investigation into the human tactile sense (see, for example, (Cholewiak and Sherrick, 1981), (Schneider, 88)), tools to support hand rehabilitation (see, for example, (Wise et al, 1990)), and devices to support teleoperation. This previous, and still ongoing, research has provided much useful information for research and development on VE tactile interfaces. Developments from the teleoperation area have been particularly useful since the prime difference between teleoperation and VE tactile interfaces lies in what drives the tactile displays: tactile sensors in a remote environment or computer models. Much of the current research on tactile interfaces discussed in the literature is presented in the context of both teleoperation and VE application.

The remainder of this section discusses the work of individual research groups who are investigating tactile displays for use in VEs. In addition, two force feedback displays under development also provide tactile feedback; these are discussed in Section 6.2.3.1 and Section 6.2.3.2. Identified efforts for which information was not available include: work on the development of a tactile display glove at the Georgia Institute of Technology; work on tactile displays under Dr. Fearing at the University of California, Berkeley; and Dr. Canepa's investigation of piezoelectric and electrorheological materials for tactile displays at the Universita di Pisa, Italy.

6.1.3.1 Armstrong Laboratory

In the Armstrong Laboratory, Crew Systems Directorate, Human Sensory Feedback for Telepresence Project, researchers led by Capt. Chris Hasser have been looking at the use of tactile feedback for telerobotic and VE applications. In one study, these researchers conducted an evaluation of the perceptual characteristics of a 5 x 6 element array tactile stimulator, with elements spaced 3 mm apart in each direction. The actuators for this device were SMA wires, used to cause the tactile elements to rise and fall. In an experiment, three subjects were tested to see if they could perceived patterns presented with the device (Hasser and Weisenberger, 1993). Two sets of stimulus patterns were used. The first set consisted of eight static patterns built of one or two straight lines. The second set consisted of the same eight patterns, presented in successive frames to simulate movement across the finger. For an eight pattern set, chance identification is 12.5% correct. The subjects gave significantly higher scores for both sets of stimulus patterns. For the static patterns, correct identifications were made 90 to 100% of the time, and 80 to 100% of the time for the dynamic patterns. In both cases, varying stimulus frequencies gave little difference in performance. Together with a physical evaluation, this experiment demonstrated that SMA arrays have the potential for presenting complex information, such as that required to represent local object shape and surface texture. The tactile feedback array has been adapted for use in presenting a virtual tactile surface to the user. This device is called the HAPtic-TACtile (HAPTAC) and is itself being used in the TacGraph system to present data plots to blind persons. Since this early work, however, the researchers have become concerned that the bandwidth of present SMA arrays may be insufficient for many haptic exploration applications. They funded a Small Business Innovation Research project to improve SMA technology and found that higher bandwidth could only be achieved at the expense of more complex, heavier apparatus.)

In more recent work, the researchers have been investigating the integration of a single element tactile stimulator with the PHANToM force feedback system (see Section 6.2.2.9). With the PHANToM, the force feedback is delivered via a thimble into which the user inserts his finger (alternatively, the forces can be delivered via a stylus that is held by the user). The tactile stimulator was required to be capable of delivering both steady-state and vibratory forces. A key concern was to use a tactile actuator capable of adequate force with a mass low enough to avoid compromising the PHANToM's dynamic performance. In

addition, to represent hard surfaces, the stiffness of the actuator system needed to be higher than that of the fingertip. The third requirement was to provide a bandwidth high enough for accurate representation of dynamic environments. After consideration of four actuator options (an electric motor with threaded screw reducer, SMA wires, pneumatic pistons, and solenoids), a solenoid actuator was selected for use. The resulting special tactile feedback thimble attaches to the PHANToM gimbal and is secured to the user's finger by means of velcro straps. Initial performance evaluation of the system found that an adaptive proportional-integral algorithm using continuously variable gain scheduling helped to compensate for nonlinearities in the solenoid actuator. The closed-loop behavior met the performance requirements of a maximum tactor force of 2 N and steady-state force accuracy of less than 0.12 N. The mass added to the force feedback system, however, degraded PHANToM's performance. Improvements in future prototypes are expected to reduce the mass of the tactile feedback hardware by over 30% (Hasser and Daniels, 1996).

Another effort with a force reflecting interface is looking at the application of Fitt's Law in VEs. For example, how scaling differences between finger movement and cursor movement impact a tapping task. Other factors, such as the addition of virtual masses to the fingertip, or viscous damping fields, may improve or degrade performance.

6.1.3.2 Begej Corporation

Begej Corporation is developing large-scale tactile displays under contract to NASA Johnson Space Center. The technology used in the tactile displays is expected to be patented and few details are currently available. What is known is that a large-area display, using 512 tactile elements, that can be worn over the upper torso, lower arms, and upper arms is being developed, together with a fingertip display using 37 tactile elements. The devices being developed may result in commercial products.

6.1.3.3 Harvard University

Dextrous manipulation for teleoperation is one of the research areas at Harvard University's Division of Applied Sciences and, for several years, researchers have been looking at tactile sensing and display devices to support such manipulation. One of the goals of this work is the development of a tactile shape display for use in the grasping surface of a force-reflecting master robot hand. Another is to determine the utility of vibration feedback and delineate the types of tasks where high frequency vibration information is important. Accordingly, these researchers, led by Dr. Robert Howe, have developed prototype tactile displays that deliver shape or vibration feedback and have conducted a series of studies using these displays.

The tactile shape display uses blunt (piano wire) pins driven by SMA actuators. Specification details for the display are given in Figure 69.

Specification	
Pin Tip	1.7 mm diameter
Pin Spacing	2.1 mm center-to-center
Rise Height	3 mm
Rise/Fall Time	62 msec
Pin Placement	4 layers of 6 pins
Pin controllability	Individually addressable
Force	1.2 N
Bandwidth	6 - 7 Hz, operates at 3dB
Display Dimensions	67 length x 26 width x 31 height (mm)

**Figure 69. Prototype Tactile Shape Display
(Harvard University)**

information in telemanipulation, two shape displays were mounted on a master manipulator. This device was intended for use in precision pinch grasp operations, with the tactile displays providing shape information to the tips of the thumb and index finger, and the master manipulator providing contact forces. (The master manipulator was a two-fingered hand with 2 DOFs in each finger, controlled using a conventional bilateral force reflection control scheme. It operated with a force reflection bandwidth greater than 80 Hz, a rise time delay for the force feedback of 15 msec, and was capable of providing 0.7 N.) The remote slave manipulator used was a two-fingered hand very similar to the master manipulator, with tactile array sensors mounted on each robot fingertip to provide shape measurement. One study already completed using this system looked at subjects' ability to localize tactile features using the device (Kontarinins and Howe, 1995). The task chosen for this study was a simulation of tumor localization using palpation. The tumor was simulated by embedding a cylindrical 4 mm diameter piece of hard rubber beneath the surface of a block of foam rubber. For the experiment, a single row of the tactile array sensor and a single row (6 pins) of the tactile display were used. Subjects performed the task both with and without the tactile shape feedback. Force feedback was provided by the master manipulator in both conditions, but visual feedback was not provided. A total of 60 trials were performed by three subjects. When the tactile feedback was available, subjects located the tumor with an error ≤ 1 mm in more than 50% of the trials, and with an error ≤ 3 mm in more than 95% of the trials. When the shape information was not available, the mean absolute error was >13 mm.

Current work on tactile shape interfaces is following two directions. In one, the researchers are looking at ways to increase the bandwidth of the display to around 25 Hz. In the other, they are looking for inexpensive ways in which to manufacture such a device. Future work is expected to focus on identifying the tactile feedback bandwidth and dynamic range requirements needed for different tasks and developing a detailed specification for system performance. Additional work will include integrating the tactile feedback system with surgical instruments such as laparoscopic forceps.

With respect to their work with high frequency vibration feedback, the researchers have developed a prototype display that uses voice coil actuators assembled from miniature 0.2 watt loudspeakers. This prototype has a 3 mm range of motion, a peak inertial force of

In an informal study on the functionality of the tactile shape display, subjects initially were asked to classify patterns generated by the display as a point, line, or plane, and subsequently to distinguish between four different orientations of lines. All subjects correctly answered all tests providing initial confirmation that the display did generate recognizable spatial patterns (Kontarinins and Howe, 1995).

In order to investigate the role of shape

0.25 N at 250 Hz, and physical dimensions of 67 x 26 x 31 (mm). As discussed by Kontari-
nis and Howe (1995), for the experiments outlined below two of the displays were mounted
on the fingertips of the master manipulator described previously, and skin acceleration sen-
sors were mounted on the slave manipulator to measure the vibrations to be produced by
the tactile displays.

Experiments have been conducted that examined the utility of this type of display
for three categories of tasks, that is, tasks where (1) the detection of vibration is the funda-
mental goal of the task, (2) vibrations indicate the state of the task, and (3) vibrations are
not directly important to the task. For the first experiment, five subjects were asked to use
touch inspection to distinguish a worn ball bearing set from a pair of such sets. Four feed-
back conditions were used: no haptic feedback, force feedback only, vibratory feedback
only, and both vibratory and force feedback. Two protocols were used: in the first, subjects
rotated both bearings in order to distinguish the worn set, and in the second they had to
make the decision based on examination of only one bearing. Eighty trials were completed
for the first protocol and 120 trials for the second. When the two set of bearings were avail-
able for examination, with no haptic feedback subjects made the correct selection in only
50% of the trials. Force feedback improved this result to 80% ($p \sim 0.1$), and with vibratory
feedback the subjects achieved 100% success with or without force feedback ($p < 0.025$).
When only one set of bearings was available for examination, with no haptic feedback the
correct response rate was 53%, with force feedback only this rose to 73% ($p \sim 0.1$), and with
vibratory feedback only the correct response rate was 66% ($p < 0.05$). With both types of
haptic feedback, and the correct response rate rose to 90% ($p < 0.025$) (the researchers note
that the subjects had difficulty in manipulating the bearing in the time provided without
force feedback).

In second experiment subjects used the master manipulator to control the slave
manipulator in piercing a 0.05 mm thick plastic membrane while minimizing the force
used. For this task, a sharp needle was held between the fingers of the slave manipulator.
The same feedback conditions were used as before, and three subjects performed a total of
152 trials. The force exerted during a trial was measured and used to determine subject
reaction time and any excess force exerted. The results showed that the presence of either
vibratory or force feedback significantly decreased mean reaction time by approximately
one half that obtained when no haptic feedback was provided ($0.005 < p < 0.025$). The com-
bination of vibratory and force feedback further reduced reaction time by approximately 50
msec.

In the final experiment, subjects were asked to perform a close-fit peg-in-hole
assembly task as fast as possible. Here precise control of contact forces was the critical ele-
ment and the task is an example of cases where vibrations are not directly important to the
task. While the vibration feedback did not have any significant effect on task completion
times, the researchers note that the subjects gave subjective reports indicating that the sys-
tem felt more with "complete" when the vibration display was used.

Currently, the researchers are looking at medical applications for vibratory feedback, and mounting a tactile vibration interface system on such tools as catheters and biopsy needles. As a separate effort, commercialization opportunities for this technology are being investigated.


6.1.3.4 Hokkaido University, Japan

Led by Dr. Shuichi Ino, for the last several years researchers at Hokkaido University have been investigating the development of an integrated system of displays for providing sensory feedback to a human hand. A large part of this work has concerned the development of tactile displays for presenting shearing and pressure forces, and for presenting temperature feedback.

The researchers have experimentally examined the human capability for passive perception of shear, as reported in Section 6.1.1. The results of this experimentation yielded tactile display design requirements in the areas of strain generation mechanisms and temperature, and a test production device capable of generating 3-D micro displacement of shearing and pressure sensations has been developed. This device uses a pneumatic system to separate the display device from the driving mechanism, enabling a small and lightweight (22.5 g) display. The air pressure on each cylinder is computer-controlled using an electro-pneumatic regulator, controlling both the pressure and shearing sensations generated by means of a lateral-moving stage. The stage stroke is ± 3 mm on both x and y axes. The maximum pressure output is 600 gf. Current work with the display is focusing on the development of a device suitable for mounting on a fingertip and further evaluation of its psychophysical characteristics.

With respect to temperature feedback, the researchers have conducted experiments that investigated human ability to recognize different materials (aluminum, glass, rubber, polyacrylate, and wood) based on differences in fingertip skin temperature when touching the material (Ino, 1993). The distinguishing factor was found to be temporal temperature difference. Using this information, a tactile temperature display was developed. The temperature of the display surface is measured by a thermocouple and a Peltier module allows the display to act as both a heater and cooler. A photograph and further details for this display are given in Figure 70.

Using this device, and the temperature change patterns acquired in the psychophysical experiments, an experiment was conducted to assess the effectiveness of the display in allowing users to distinguish between objects based on temperature feedback. In this experiment, artificial thermal stimuli were presented to four subjects who were asked to identify the material. Analysis of the results showed no significant difference between identifications made using the real materials and the temperature display. However, neither form of identification was completely correct in every case, and the researchers suggest that the presentation of temperature information be used as just one element of tactile feedback systems intended to support absolute material recognition.



Specification
 Temperature Range -10 to 60° C
 Temperature Resolution 0.1° C
 Display Dimension 16 sq cm
 Display Weight 30 g
 Interface RS-232

Figure 70. Temperature Display (Hokkaido University)

6.1.3.5 Hull University, UK

Researchers at Hull University, Department of Electrical Engineering, are looking at the use of electrorheological fluids for tactile displays. The fluids under consideration are primarily a colloidal dispersion consisting of an insulative base oil and a slightly conductive dielectric solid particulate. Under the stimulus of an electric field, these fluids have the ability to change from a liquid to a pseudo-solid state almost instantaneously and their malleability is dependent on the strength of the electric field. One of the advantages of this type of display is the absence of moving parts, if only the display of contact and shape information is required. The presentation of surface textures would require a system of control electronics.

The researchers, headed by Dr. Taylor, have developed a single cell display and are now working to develop a second generation display that will employ an array of electrorheological elements. This display is expected to be ready to embed in a VE tactile interface device within the next year.

6.1.3.6 Massachusetts Institute of Technology

At the Massachusetts Institute of Technology (MIT), Department of Mechanical Engineering, researchers in the Touch Laboratory are looking at human haptics and its relationship to machine haptics. This work is being led by Dr. Mandayam Srinivasan. The overall goals are to (1) develop an understanding of the human as the perceiver of, and the operator on, the environment; and (2) to apply this basic knowledge in the areas of rehabilitation, robotics, and human-machine interfaces for VEs and teleoperation. In particular, current work is focusing on haptic information acquisition and the control of contact tasks with the hand, with an emphasis on the associated information processing mechanisms. It includes investigation of the biomechanics, neurophysiology, and psychophysics of touch, and the development of a computational theory of haptics. Collaborators in much of the

work include Dr. Nat Durlach and Dr. Ken Salisbury from MIT, and Dr. Robert LaMotte from the Yale University School of Medicine.

As part of a project focused on haptic interface development for VEs, Dr. Srinivasan's group has developed two major devices for performing psychophysical experiments, the Linear and Planar Graspers. These are in use, along with the PHANToM (see Section 6.2.2.9). Software to allow the haptic devices to present simulations of fundamental mechanical object properties, such as compliance, viscosity, and mass; to display shape, texture, and friction of solid objects; and to portray virtual walls and corners has been developed. Initial psychophysical experiments have measured the manual resolution of stiffness, viscosity and mass, investigated the influence of visual information on haptic perceptions of stiffness, and looked at the feasibility of various haptic display algorithms for presenting the shape, texture, and friction of solid surfaces. These experiments have yielded insights that show how human sensory perceptions can be used to promote haptic sensations in the user. For example, one finding is that the perception of the stiffness of objects like virtual push-buttons can be significantly altered by presenting visually skewed positional information to the subject. Additional psychophysical experiments are underway, aimed at characterizing the effectiveness of refined, computationally-efficient simulations and rendering algorithms in conveying desired object properties to the human user.

Recently, a haptic rendering technique called "force shading" (analogous to "Phong shading" in graphics) has been demonstrated to give users the feel of smoothly curved surfaces, even when the surfaces are represented as polyhedrons. In investigating multimodal displays, the effect of contact sounds on the perception of object rigidity is being explored.

The development of hardware and software haptic interfaces for human interactions with multimodal VEs will be continued in future work. This work is expected to include the development of high performance tactile sensors and displays, as well as a variety of haptic rendering algorithms that take advantage of human illusions in perceiving multimodal sensory inputs.

6.1.3.7 Research Center at Karlsruhe, Germany

Researchers at Karlsruhe Research Center, Department of Engineering Technology, are developing a tactile feedback system for use with flexible endoscopic forceps. These researchers, led by Dr. Harald Fischer, have developed a tactile display that consists of three 24-needle printing heads thus providing a total of 72 actuators, although only 64 are actually used. Individual needles are electromagnetically triggered by opto-decoupled printout boards and vibrate at a maximum frequency of 600 Hz to present contact pressure sensations. The tactile display is mounted on a box. It is driven to respond to operator applied forces detected by a force-movement sensor placed in the distal shaft of the forceps, grasping forces applied to the tissue that are detected by a miniature pressure transducer, as well as pressure distribution between tong and tissue as measured by a tactile sensor placed between the jaws of the forceps. In this way, the distribution of pressures and handling forc-

es are sensed and displayed to the fingertip of the surgeon, they are also displayed graphically on a PC screen and sent to a plotter. A technical specification of this device is not publicly available at this time.

Future research on the tactile interface system will focus on the development of an analog linear device with 64 needles that will be mounted directly at the end of the laparoscopic forceps so that the surgeon can operate the system with a single finger. An optical sensor array for the distal end of the forceps will also be developed.

6.1.3.8 Sandia National Laboratories

Dr. Dave Andaleon at Sandia National Laboratories, is leading researchers in developing fingertip tactile feedback devices for VE applications. Specifically, the goal of this work is the development of a high density tactile array compatible with standard VE device interfaces.

After a review of haptic feedback research and products, the researchers developed a set of quantitative and qualitative evaluation metrics and a tactile feedback testbed for evaluating tactile stimulus technologies (hydraulic/pneumatic, electro-magnetic, piezoelectric, and bi-metallics such as shape memory alloys, polymeric gels, electrorheological, and magnetostrictive materials). On the basis of these evaluations, it was decided to build a electromagnetic actuator. (A patent for the actuator design is pending.) This actuator operates in the frequency range 8 - 100 Hz, it is capable of 762 micron indentation and exerting a maximum pressure of 1.2 N/cm^2 .

For the tactile display itself, a 2×3 array of actuators is mounted on a pad, and pads are attached to a user's fingers using velcro straps. The software developed to support the tactile interface system allows tactile displays to be used on the thumb, index finger, middle finger, and palm simultaneously. Each actuator in a tactile display is individually controlled with respect to magnitude, frequency, and phase. A serial RS-232 interface is provided through a host computer with analog output boards. A performance specification for the tactile display, and a photograph, are given in Figure 71.




Figure 71. Tactile Display (Sandia National Laboratories)

Initial tests with the tactile interface system included the simple mapping of material texture to actuator frequency and modeling cubes as solids using object collision detection. The stimulus types investigated were magnitude, frequency and phase, spatial and temporal frequency, and spatial and temporal patterns.

Current work is focusing on developing software that uses the tactile display to present a variety of textures and other surface information. The resulting tactile interface system will be integrated into a situational training VE that supports multiple participants. Insights gained from this use of the tactile interface system will be used in further investigation of actuator performance and actuator ruggedness. Future work also is expected to look at the value of providing tactile feedback in the absence of any kinesthetic feedback.

6.1.3.9 TiNi Alloy Company

Under contract with the Human Systems Center at Brooks Air Force Base, TiNi Alloy has developed a tactile display consisting of a 5 x 6 array of tactor pins. A photograph and details for this display are given in Figure 72. The display is supported by microcontroller hardware and software to constitute a complete tactile system.

An initial informal study of the effectiveness of the tactile display has been performed. In this study, when the tactile display was mounted on a digitizing puck, so that the user's fingertip rested on the tactile display while his hand moved the puck across a flat surface, subjects were able to correlate patterns shown on a screen and those presented via the pins. In a more formal, but preliminary evaluation, three subjects were able to identify a set of static patterns and a set of moving patterns (Hasser and Wesenberger, 1993). Currently, the display is being refined to support its efficient manufacturing. As part of the same contract effort, TiNi Alloy's engineers are augmenting the force feedback provided by the PHANTOM with tactile feedback; this is expected to be achieved by mounting a single tactor in the PHANTOM thimble and activating this tactor remotely, perhaps using pneumatic actuators.




Figure 72. Programmable Tactile Array (TiNi Alloy)

In an effort funded by the Naval Sea Systems Command, TiNi Alloy is further investigating the use of its SMA thin film pneumatic micro-actuators for actuating tactile feedback. The intended application of this work is a tactile display that can be positioned on a pilot's torso and used to alert him to special circumstances.

6.1.3.10 University of Salford, UK

A glove with tactile, contact pressure, and temperature feedback, referred to as tele-taction, is being developed by researchers at the University of Salford, Department of Electronic Engineering.

A tactile sensation of texture and slip is provided for object identification and grasp stability control using vibrational stimulation from a piezo-electric actuator. The feedback module is a PZT (lead zirconate titanate) ceramic disc, 10 mm in diameter and 1 mm thick, that is mounted on a metal disc 15 mm in diameter and 1 mm thick. This transducer is enclosed in a PVC film and driven by a high voltage (up to 350 V). The total unit weight is around 2 g. Finger positions are sensed using Hall Effect sensors that provide measurement of the metacarpophalangeal (MCP) and proximal interphalangeal (PIP) finger joints. These various devices are mounted on a glove to provide feedback to a single finger. Contact forces, or pressure, are provided by pneumatic bladders that are operated by an independent pneumatic powerpack, connected to the glove via valves and piping. This feedback is transmitted to thirty locations on the anterior surfaces of the fingers and palm of a hand. In tests of the texture feedback, subjects were able to distinguish between ribbon cable, writing paper, tissue paper, a small file, cloth, and four different textured steel plate surfaces. They also were able to detect slips of 0.5 mm or more. In the case of pressure feedback, the subjects were able to distinguish between four different force levels (2, 10, 30, and 60 N).

Thermal feedback using a Peltier Effect Heat Pump supports object or material identification and, also, safety. A rapid response thermal-couple is mounted on the Peltier unit, in contact with the user's skin, allowing tracking of the user's skin temperature and the provision of a rate of cooling or heating relative to this temperature. A small aluminum plate is attached to the exterior surface of the heat pump to act as a thermal regulator that minimizes the temperature gradient and a small heat sink with an integral fan unit permits high cycle rate responses. This thermal device is set to generate temperatures in the -5 to 50°C range, with rate changes of up to 20°/sec. With thermal feedback, subjects were able to distinguish between five objects with different temperatures or thermal conductivities (ice cube, a soldering iron, insulating foam, aluminum block, and room condition). Finally, temperature was used as a substitute for pain or danger feedback by rapidly increasing the temperature to 50° C; in tests, subjects were able to respond to this feedback with a reaction time of 0.9 sec. The feedback unit weighs 10 g, measures 15 x 15 x 3 mm and operates at 10 W.




Figure 73. Tactile Feedback Glove
(University of Salford)

include sensors for the detection, at a minimum, of pressure, vibration, and temperature and so provide data on such features as object shape, profile, and hardness. The sensors need not be limited to normal human sensations and may be used to provide information on characteristics such as conductivity and radioactivity, encoded as tactile feedback.

With respect to VEs, the researchers' objective is to create an effective feel for virtual objects ranging from switches and levers to walls. They plan to use the sensor inputs acquired from robot contact with real objects to program corresponding features for use in VEs.

6.2 Kinesthetic Interfaces

Since the majority of a human's interaction with his environment consists of manipulating objects, this capability is a prerequisite for many practically useful VEs. Manipulation can be achieved via indirect means such as voice, keyboard, or mouse commands. A natural form of manual manipulation, however, requires use of the types of interaction devices discussed in Section 5, augmented with force feedback to simulate object properties such as overall shape, stiffness, and weight. The importance of force feedback has been well established in the teleoperation community (see, for example, (Hill and Salisbury, 1976), (Hannaford, 1989), and (Howe, 1992)). There have also been some experiments that have investigated the value of force feedback for VEs. In a molecular docking VE, Ouh-Young, Beard, and Brooks (1989) demonstrated how a visual interface supported by force feedback gave significantly better performance than the visual display alone and, when only a single type of display was available, force feedback gave a better task performance than the visual display. More recently, it has been shown that force feedback for simple grasping tasks can reduce task error rate and learning time by over 50% (Gomez, Burdea, and Langrana, 1995).

Essentially, there are three components to providing a force feedback interface for VEs: measurement of the movement of the user's fingers, hand, and/or arm, and sensing any forces he exerts; calculation of the effect of the exerted forces on objects in the VE and the

resultant forces that should act on the user; and presentation of these resultant forces to the user's fingers, wrist, and arm as appropriate.

Force feedback devices are either earth-grounded, off-the-body devices or exoskeleton devices worn by the user that, themselves, are either anchored to the ground or a body part closer to the point of force application. Hasser (1995) discusses this basic difference and also provides a discussion of the different types of actuators and transmission methods used for force feedback devices. A summary of some of the actuator technologies, taken from Hasser's report, is provided in Table 15. Of these technologies, electromagnetic motors, hydraulics, and pneumatics are technologies in current use. Piezoelectric and magnetorestrictive technologies are still the subject of research and development.

There are some commercial VE systems that have limited force feedback capability. Currently, these are all entertainment systems where the user "operates" some vehicle and forces are presented to the user via some control device such as a steering wheel. Within the next several months, however, some surgical training systems that provide force feedback via surgical instruments are expected to come to market. The best example of a non-commercial VE system in practical use that employs force feedback is the molecular docking system at the University of North Carolina, see Section 6.2.3.12 below.

6.2.1 The Human Kinesthetic Sense

As discussed by Boff, Kaufman, and Thomas (1978), kinesthesia provides humans with an awareness of the position and movement of body parts, whether such movement is self generated or externally imposed. The receptors that support this sense are found in skin, joints, and muscles. The relevant skin receptors provide information about skin stretch and cutaneous deformation and were discussed previously. Joints contain two types of receptors: Golgi endings found in joint ligaments, and Ruffini type endings found in joint capsules. These receptors respond to joint torque and capsule stretch, respectively. They are slowly adapting and thought to signal extremities of joint flexion and extension. Muscles also contain two types of receptors, Golgi tendon organs that monitor muscle tension, and muscle spindle organs that measure muscle stretch and its rate of change.

Together, these various receptors provide information about joint angles, muscle length and tension, and their rates of change. However, the most important receptors for kinesthesia seem to be the muscle spindle organs. These receptors are thought to be the primary candidate for static position detection and, probably with skin receptors, they provide a sense of movement. But none of the skin, joint, or muscle receptors provide awareness of weight or effort; instead, this sense seems to arise mainly from signals derived entirely within the central nervous system.

It is important to note the asymmetric nature of the human somatosensory system, that is, the fact that the force control and perceptual bandwidths of the human differ. For example, Brooks (1990) reports that the maximum frequency with which a typical hand can

Table 15. Force Feedback Actuator Technologies^a

Technology	Description	Advantages	Disadvantages
Electromagnetic Motors	Electromagnetic motors produce torque with two time-varying magnetic fields, caused by two coils or a coil and a magnet	- Easy to control - Clean, quiet - Easy design and installation	- Heavy components - Low power densities at small scales - Heat dissipation problems - Low static force capability
Hydraulics	A hydraulic fluid is pressurized by a power plant, controlled by servo-valves and delivered to rotary or linear actuators through pressurized fluid lines	- Force capability, power output, stiffness, and bandwidth unmatched by other technologies	- High mass - Tendency for fluid leaks - Design difficulty - Expensive
Pneumatics	A gas (normally air) is pressurized by a power plant, controlled by servo-valves, and delivered to rotary or linear actuators through pressurized fluid (air) lines	- Good static force capability - Lighter than hydraulics - Pneumatic power plants and distribution systems easier to manage than hydraulics	- Relatively low bandwidth - Low actuation stiffness - Low power capability
Piezoelectric	Piezoelectric motors translate the vibration of piezoelectric materials to linear or rotary motion using frictional forces to produce usable torques or forces at low speeds, without the need for gear reduction.	- High forces at low speeds in small package	- Requires precision machining - Necessary power gating can cause annoying and potentially hazardous noise, depending on the design
Magnetorestrictive	Magnetorestrictive materials change shape when subjected to magnetic fields. Magnetorestrictive motors also mechanically rectify small oscillatory motions of the driving element(s).	- High forces at low speeds in small package	- Necessary power gating can cause annoying and potentially hazardous noise, depending on the design - Heat dissipation can be a problem - Requires precision machining
Shape Memory Alloy	SMA wires and springs contract when heated and expand again as they cool under stress.	- Good power-to-mass ratio	- Low efficiency during contraction - Heat dissipation problems limit relaxation rate of wires - Limited bandwidth

a. Based on Hasser (1995, 1996)

transmit motion commands to a hand master is 5 - 10 Hz, while the upper bound for receiving position and force signals is not less than 20 - 30 Hz.

Researchers at various research laboratories and departments at MIT are collaborating in experiments to collect human factors data (Tan et al, 1994). This data will be used to develop a detailed catalog of human factors data that aids better design and evaluation of haptic interfaces. These researchers report that the JND for force sensing is around 7%, regardless of reference force or body site. The force required for a human to perceive an object as rigid ranges from 153 to 415 N/cm. The maximum controllable force ranges from 16.5 to 192.3 N, increasing from the most distal finger joint to the shoulder joint (here there are significant gender differences). Force output resolution is about 0.36 N regardless of body site, while, in terms of percentages, the resolution tends to decrease from the PIP finger joint to the shoulder joint. The JND for pressure perception is roughly 0.06 - 0.09 N/cm, regardless of contact position. The greater the contact area, the more sensitive the human arm is to pressure; the JND decreases by a factor of roughly four (from 15.6% to 3.7%) when the contact area increases by a factor of sixteen (from 0.2 in² to 3.14 in²).

On the whole, these figures are consistent with the findings of other researchers. For example, Shimoga (1993a) reports that the human fingers can sense force variations of 0.5 N; if this load is distributed, the pressure must not be below 0.2 N/cm² which is the minimum pressure that a human finger can sense. In summarizing much available data, Shimoga also states that human fingers can exert 30 - 50 N for brief periods, and 4 - 7 N for sustained periods. Massie and Salisbury (1994) have found that, in practice, a virtual surface with a stiffness of at least 20 N/cm is perceived as a solid and immovable wall by users.

Table 16. Variability of Forces Exerted in Human Grasping

						
5% Female	53 lbs	53 lbs	7.5 lbs ⁷¹	7.5 lbs	9 lbs	4 lbs ⁷³
95% Male	147 lbs	147 lbs	30 lbs ⁷¹	30 lbs	32 lbs	13 lbs ⁷¹
Torque Capability	Excellent	Excellent	Good	Poor	Some	Excellent
Endurance @25% Load	Good	Good	Poor	Fair	Fair	Good

7. Data unavailable; 1. Values assumed to be about the same as pincher grasp but supporting evidence not available; 2. Mean value 100 male subjects; 3. Value assumed to be 1/3 of male value.

Reprinted by permission of Taylor & Francis Group Ltd.

Jacobus et al (1992) have summarized the variability of force output in different types of grasp as shown in Table 16. As indicated in this table, humans use several different types of hand grasp in manipulating objects, and the functional characteristics of these grasps differ. Since the current technology uses different methods for

providing force feedback depending on the type of grasp used, it is useful to briefly delineate these different types. Schlesinger (1919) first categorized grasps as cylindrical, finger-tip, hook, palmar, spherical, and lateral. Since the type of grasp used tends to reflect the task to be performed, Napier (1956) suggested a categorization based on the distinction between power grasps and precision grasps. Additional schemes have been based on the concept of virtual fingers, oppositions provided by various hand configurations, and in terms of prehensile and non-prehensile grasps. Cutkosky and Howe (1990) provide a grasp taxonomy that relates these different categorization schemes, as shown in Figure 74. They also define grasp attributes as dexterity, precision, sensitivity, stability, and security.




Figure 74. Taxonomy of Manufacturing Grasps

Reprinted by permission of the Robotics Institute, Carnegie-Mellon University.

Grasping does not rely only on kinesthetic sensing, tactile signals play a significant role in almost all manipulative tasks. Johansson (1991) argues that cutaneous surface deformations directly reflect the accomplishments of many manipulative actions, and may serve as preconditions for triggering some of the motor commands associated with these actions; tactile signals also provide information about an object's physical properties that are used in guiding the use of manipulative forces.

Experiments conducted in a teleoperator environment provide some data on operator fatigue and discomfort that might arise using hand force feedback interfaces in a VE, and that can reduce an operator's ability to estimate force magnitudes and variations (Wiker, 1989). The factors that aggravate fatigue and discomfort are cited as grasping force and work-to-rest ratio. Operator comfort will be within safe levels if the grasping, or reflected, forces are less than about 15% of the maximum exertable force, that is, the index, middle and the ring fingers can safely exert about 7, 6, and 4.5 N, respectively, without encountered fatigue and discomfort.

6.2.2 Commercially Available Devices

This discussion is limited to force feedback devices that are specifically intended for use in VEs. Even so, recent years have seen several devices, of quite different types and capabilities, come to market. The features of these devices are summarized in Table 17. Information on another commercial product, Sarcos Inc.'s Hand Master, was not available.

In addition to these existing products, Virtual Technologies, Inc. is currently developing its *CyberForce* product that will augment the *CyberGlove* with force feedback. This new force feedback device will provide restrictive forces to the user's fingertips and is expected to be released in Summer '96.

6.2.2.1 4 DOF Force Feedback Master (Surgical Simulator)

Initially designed for use in medical simulations, the EXOS, Inc. 4 DOF Force Feedback Master provides force feedback to the hand and arm in 4 DOFs. Feedback is provided via a handle connected to a larger tool shaft, which is pivoted at one point with an active 3 DOF gimbal. The fourth DOF is provided by a linear sliding module, allowing the tool to be translated ("heaved") along the shaft of the tool. One possible application of the device is the simulation of minimally invasive surgery, in which forces encountered by touching virtual tissue and organs with a laparoscopic tool are simulated and displayed to the user at the handle. The additional information available for this device is given in Figure 75. The 4 DOF Force Feedback Master is made to order and pricing information is not available.

6.2.2.2 Force Exoskeleton ArmMaster


Structured as an exoskeleton, the EXOS, Inc. Exoskeleton Force ArmMaster (EAM II) has 5 active DOFs and additional passive freedoms designed for comfort and the ability to adjust to different arm sizes. An active gimbal structure suspended above the

Table 17. Characteristics of Commercially Available Force Feedback Devices

Device Name	Vendor	Device Type	Forces Provided To	Active DOFs	Force Resolution	Applied Force ^a	Price
4 DOF Force Feedback Master	EXOS, Inc.	Desktop	Hand via joystick	4	Unavailable	Range 5.1 - 12.0 oz-in (cont), range 20-59 oz-in (peak)	Contact vendor
Force Exoskeleton ArmMaster	EXOS, Inc.	Exoskeleton	Shoulder and elbow	5	Unavailable	Range 3.4 - 56.6 in-lb (cont), range 29.0-489.0 in-lb (peak)	Contact vendor
Impulse Engine 3000	Immersion Corporation	Desktop	Hand via joystick	3	0.00435 N	8.9 N (cont)	\$7,950
Laparoscopic Impulse Engine	Immersion Corporation	Desktop	Hand via tool handle	5	0.00435 N	8.9 N (cont)	\$8,950
Interactor	Aura Systems, Inc.	Vest	Torso via vest	N/A	N/A	N/A	\$99
Interactor Cushion	Aura Systems, Inc.	Cushion	Back via cushion	N/A	N/A	N/A	\$99
HapticMaster	Nissho Electronics Corporation	Desktop	Hand via knob	6	2.85 gf	1.2 kgf, 5.6 kgf/cm (cont), 1.8 kgf (peak)	Unavailable
Hand Exoskeleton Haptic Display (HEHD)	EXOS, Inc.	Exoskeleton	Thumb & index finger joints, palm ^b	4	Unavailable	Range 1 - 5 lb (peak)	Contact vendor
PER-Force 3DOF	Cybernet Systems Corporation	Desktop	Hand via joystick	3	0.035 oz	1 lb (cont), 9 lb (peak)	\$9,950
PER-Force Handcontroller	Cybernet Systems Corporation	Desktop	Hand via joystick	6	12 bit	2- 3 oz (min), 20 - 25 lb (peak)	Contact vendor
PHANToM	SensAble Devices, Inc.	Desktop	Fingertip via thimble	3	12 bit	1.5 N (cont), 10 N (peak)	\$24,000
SAFIRE	EXOS, Inc.	Exoskeleton	Wrist, thumb & index finger	8	Unavailable	Range 1 - 2 lb (peak), wrist 2 lb (peak)	Contact vendor

a. Some figures are for torque, not force

b. Includes tactile display presenting a sense of slip to the thumb and index finger.




Specification			
Force/Torque		Cont.	Peak
	Pitch/yaw	12 oz.-in	59 oz-in
	Roll	5.1 oz-in	20 oz-in
	Heave	7.5 oz	25.3 oz
Resolution	Pitch/yaw	0.1°	
	Roll	0.1°	
	Heave	0.001 in	
Range of Motion	Pitch/yaw	140°	
	Roll	350°	
	Heave	4 in	
Device Size	12L x 13W x 7H (in)		
Device Weight	~6 lb		
Interface	Serial, custom		

Photo courtesy of EXOS, Inc.

Figure 75. 4 DOF Force Feedback Master

shoulder provides 3 DOF force feedback to the upper arm. A remote center mechanism provides 2 DOF force feedback to the lower arm. The active DOFs on the shoulder use DC motors with a closed loop cooling system that allows the motors to produce twice the usual torque without overheating. The EAM II is mounted to the arm via an air bladder that accommodates small misalignments of the device. The system is completely back-mounted and designed to be lightweight and portable. If desired, position sensing, using optical encoders, provides motion commands to the simulation or slave. A photograph of the EAM II and further details are given in Figure 76.



Specification			
Torque		Cont.	Peak
	Shoulder Ab/Ad, F/E	56.6 in. lb.	489.0 in. lb.
	Shoulder I/E	20.3 in. lb.	175.0 in. lb.
	Elbow F/E	14.0 in. lb.	121.0 in. lb.
	Forearm S/P	3.4 in. lb.	29.0 in. lb.
Powered DOF	Shoulder	3	
	Elbow	1	
	Forearm S/P	1	
Joint Motions	Shoulder Flexion/Extension		120°
	Shoulder Ab/Adduction		120°
	Shoulder Int/External Rotation		100°
	Elbow Flexion/Extension		100°
	Forearm Supination/Pronation		100°
Backlash	<2 °		
Friction	Approx. 4% of torque		
Size	Adjusts to fit most male, female arms		
Weight	~4 lbs. on arm		
	~18 lbs. back mounted (single arm)		
	~20 lbs. back mounted (double arms)		


Photo courtesy of EXOS, Inc.

Figure 76. Force Exoskeleton ArmMaster

The Force ArmMaster can be configured for one or two arm operation, or integrated with SAFiRE or HEHD to provide force feedback to the wrist and fingers. The Force ArmMaster is made to order and pricing information is not available.

6.2.2.3 Impulse Engine Family

Based on its Impulse Engine, Immersion Corporation markets a range of tool-based force feedback devices. All Impulse Engine products use servo-motor actuators. They come with device drivers for a variety of machines ranging from PCs and Macs, to Silicon Graphics platforms. A small number of demonstration programs also are available. The Impulse Engine 3000 is a 3 DOF pen-based device priced at \$7,995; a photograph and specification details are given in Figure 77.



Specification	
Continuous Force	8.9 N
Force Resolution	0.00435 N
Position Resolution	0.01 mm
Backdrive Friction	<0.14 N
Bandwidth	650 Hz
Workspace	13 cm (linear)

Photo courtesy of Immersion Corporation

Figure 77. Impulse Engine 3000

The Laparoscopic Impulse Engine is an interface device specifically designed for virtual simulations of laparoscopic and endoscopic surgical procedures. This device can be fitted with a selection of instrumented surgical tools, or tool handles. The surgical tool can pivot (with 2 DOF) around the insertion point with an approximate range of 100° and a maximum torque of 60 oz/in. A third DOF allows for translation in-and-out along the insertion axis with a maximum travel of 4 in, and forces of up to 2 lbs. The fourth DOF allows the instrument to spin a full 360° along its longitudinal axis, and the fifth DOF provides for the open-close motion of the instrument grip. Position sensing is provided for all 5 DOFs, and force feedback for the tool pivoting and travel along the insertion axis. This device is priced at \$8,950, or \$15,950 for a pair of devices. A photograph and specification details for the Laparoscopic Impulse Engine are given in Figure 78.

Another tool-based force feedback device marketed by Immersion Corporation is the Needle Insertion Simulator. This device is intended for use in a training system that tracks the insertion of a virtual needle while providing force feedback that simulates the needle's penetration through various layers of tissue. It is composed of a single linear axis with a travel of 13 cm, and provides forces up to 8.9 N. A final tool-based device is the Virtual Catheter Interface.

Immersion Corporation is continuing refinement of its current products to develop higher performance versions. A future area of research is expected to be the development




Figure 78. Laparoscopic Impulse Engine

of virtual fixtures, that is, abstract perceptual information that can be overlaid on a virtual workspace to aid in task performance.

6.2.2.4 Interactor and Interactor Cushion

The Interactor products from Aura Systems, Inc. are very different from the other force-feedback products discussed here: they monitor an audio signal and use Aura's patented electromagnetic actuator technology to convert bass sound waves into vibrations that can represent such actions as a punch or kick. Both the Interactor vest and the Interactor Cushion plug into the audio output of a stereo, TV, or VCR. The user is provided with controls that allow adjusting the intensity of vibration and filtering out of high frequency sounds. The audio signal itself is reproduced through a speaker embedded in the vest or cushion.

The Interactor Vest is worn over the upper torso and costs \$99, further details are given in Figure 79. The Interactor Cushion is placed against a seat back and the user leans against it, its price is \$99. Further details for the Interactor Cushion are given in Figure 80.

6.2.2.5 HapticMaster


The HapticMaster was developed by Dr. Hiroo Iwata at the University of Tsukuba, Japan, and is now marketed by Nissho Electronics Corporation. This device is a desktop instrument that provides 3-D force and 3-D torque to the user via a knob grasped by the user's fingers. The actuators are three sets of pantograph linkages, each driven by three electric motors. The top of each pantograph is connected to a vertex of a small platform by



Specification
 Frequency Range 0 - 100 Hz
 Dimensions ~19 H x 13 W x 4.5 D in
 Weight 8.5 lb

Photo courtesy of Aura Systems, Inc.

Figure 79. Interactor



Specification
 Frequency Range 0 - 100 Hz
 Dimensions ~17 L x 13 W x 4.5 D in
 Weight 4 lb

Photo courtesy of Aura Systems, Inc.


Figure 80. Interactor Cushion

a spherical joint, and the knob is mounted in the center of this platform. A specification for the HapticMaster and a photograph of the device are given in Figure 81.

Software that computes positions and forces is available for the PC. The HapticMaster itself is controlled by an interface unit that provides signal amplification and A/D converters for measuring master angles.

6.2.2.6 Hand Exoskeleton Haptic Display

The EXOS, Inc. Hand Exoskeleton Haptic Display (HEHD) is an integrated multi-modal haptic display system that provides force feedback as well as a sense of slip to the thumb and index finger. The device consists of a hand exoskeleton (a modified SAFiRE, see Section 6.2.2.10) providing 1 DOF force feedback to the thumb and 2 DOF force feed-




Specification	
Peak Force	1.8 kgf
Continuous Force	1.2 kgf
Continuous Torque	5.6 kgf/cm
Force Resolution	2.85 gf
Position Resolution	0.4 mm
Backdrive Friction	0.1 kgf
Max. Payload	2.5 kg
Update rate	90 Hz
Workspace	Sphere 40 cm diameter
Interface	RS-232

Photo courtesy of Nissho Electronic Corporation

Figure 81. HapticMaster

back to the index finger. The slip displays each provide a sense of slip in one direction and are integrated into the exoskeleton. The exoskeleton can be mounted to a boom that provides 2 DOF position sensing in a vertical plane as well as force feedback in the vertical direction. When fully integrated, the system can be used for virtual pick-and-place tasks in which weight, contact, and slip information is passed by force and slip feedback.

The software that controls the HEHD is available for 386 or higher IBM-compatible PCs and the Silicon Graphics Indigo2. A photograph and specification details are given in Figure 82. Like all EXOS, Inc. force feedback products, the HEHD is built to order and general price information is not available.



Specification	
Maximum Slip display	1 lb
Force Horiz. force reflection (thumb)	1 lb
Horiz. force reflection (index)	1 lb
Vertical force reflection	5 lb
DOF Powered: Force	4
	2
	Slip 2
Passive	
Range of Thumb	120 °
Motion Index finger MCP joint	120 °
Index finger PIP joint	90 °
Vertical range of motion	7 in.
Size	Adjustable to fit most hands
Interface	VMEbus

Photo courtesy of EXOS, Inc.

Figure 82. Hand Exoskeleton Haptic Display (HEHD)

6.2.2.7 PER-Force 3DOF

Cybernet Systems Corporation markets a 3-D force-feedback, backdrivable joystick called PER-Force 3DOF. This device is primarily intended for use in VE and teleop-

eration applications. The user can move the joystick handle in three revolute directions and these are mapped to movements in terms of x , y , and z axes or angular movements in the VE. Unlike many force-feedback devices, PER-Force can be operated in different control modes. That is, the computer can read either the joystick joint or the transformed position, velocity, or force. Similarly, via small, brushless DC servo motors on each of the revolute axes, force-feedback can be presented to the user in terms of position, rate, or force. Three cueing buttons, an analog trigger, and a palm-actuated deadman safety switch are mounted on the handle; these controls are all programmable and can be used, for example, to switch device mapping between Cartesian coordinates and angular movements. A photograph and additional details are given in Figure 83.





Figure 83. PER-Force 3DOF

PER-Force 3DOF comes with a complete MS-DOS C development environment to support control system modifications and the development of custom interface drivers. The price of PER-Force 3DOF with PC-based controller is \$9,995.

6.2.2.8 PER-Force Handcontroller and Finger Forcer Option

Originally designed for use in the Space Station, Cybernet Systems Corporation's PER-Force Handcontroller is a small device that provides the user with a motorized handle with which to position robots, or virtual objects, and through which 6 DOF force feedback is provided. The handle resembles an aircraft-type sidearm-control grip with three cueing buttons, an analog trigger, and a palm-activated deadman switch. Six brushless DC servo motors are used to provide force feedback on each of the 6 axes, although lower cost versions of the device are available for 2 to 5 axis operation. The handcontroller can be operated in various modes: force-position scaling, position-position lock, rate-position orientation lock, and user-programmed axis lock. A photograph and further details are given in Figure 84.



Specification

Output Force	Max 20 - 25 lb, min 2 - 3 oz
Force Resolution	12 bits
Range of Motion	Three linear axes, 4 in; three revolute axes yaw, pitch, and roll 90°
Position Resolution	0.0003 in each linear axis, 1/90° each revolute axis
Backdrive Friction	5 oz on linear axis
Handcontroller Weight	25 lb
Handcontroller Size	14.8H x 14.8W x 10.8D in
Interface	RS-232/422, Ethernet, Parallel I/O, A-D/D-A, custom

Photo courtesy of Cybernet Systems Corporation

Figure 84. PER-Force Handcontroller

The PER-Force Handcontroller is controlled by the PER-Force Universal Robot Motion Controller, itself a 486-based PC, that provides an interface to any MS-DOS/Windows, VME, Mac, or Unix-based machine. Two development libraries come with the Handcontroller and Motion Controller, one to use in programming the controller directly, and the other for use on a host machine to which the controller is interfaced. These libraries support control system modification, reconfiguration, and interfacing; they are available in both C and X Windows formats. Additional software is provided to facilitate passing force commands to the controller. PER-Force Handcontrollers, with the PER-Force Universal Robot Motion Controllers, are custom-made and no general price information is available.

Cybernet is currently developing an additional product, the PER-Finger Forcer, that can be attached on the top of the Handcontroller to provide force feedback at the fingertips for up to four fingers and thumb. This device monitors finger position in 2 DOFs for each finger and 3 DOFs for the thumb, providing 6 DOF force feedback using miniature brushless DC servo mechanisms. The device uses thimble-like structures to grasp the user's hand, and each finger and thumb is inserted into small stirrups at the end of the effector mechanisms. It supports the full range of finger motion. The peak force output on each axis is 2 lbs, with a continuous force output capability of 0.3 lbs, and minimum force output of less than 1 oz. As with the Handcontroller, the PER-Finger Forcer supports various control modes and is driven by a PER-Force Universal Robot Motion Controller interfaced to a serial port on any MS-DOS, VME, or Unix-based computer. The PER-Finger Forcer is expected to be released on the market in Spring '96. The 3 DOF version is expected to be priced at \$9,995 and the 6 DOF version at \$59,000.

6.2.2.9 PHANToM

Developed at MIT and now marketed by SensAble Devices, Inc., the PHANToM Haptic Interface Device is a desk-based device that provides force feedback to a thimble slipped over the user's fingertip. Optical encoders are used to measure the position of the user's fingertip, with one encoder being mounted on each of 3 DC brushed motors. These motors generate forces in the x , y , and z coordinates, and the torques are passed through a pre-tensioned cable transmission to the stiff, lightweight aluminum linkage that supports the thimble. The thimble can be replaced by pen-like objects such as a stylus or scalpel to provide a tool-based interface to a VE. Specification details are provided in Figure 85. This figure also shows the initial version of PHANToM, which is priced at \$19,00. A larger version of the device, called PHANToM 1.5, with a 300% larger working space is also available for \$24,00. Even larger versions that support a full-arm workspace are available to selected research groups.




Figure 85. PHANToM

PHANToM is controlled by Silicon Graphics or 486 PC-based software and comes with a portable library of demonstrations that show how PHANToM can be used. This library is also a source of software models for various virtual objects (such as cubes, spheres, walls, and polygonally rendered objects), and additional models used for providing object properties such as texture and friction.

Two PHANToMs can be used simultaneously to support force feedback for a thumb and finger on one hand, or one finger on each hand. Additional PHANToMs can be used to provide, for example, force feedback for two fingers on each hand. All that is needed to support the use of multiple devices is special driver software that is available from SensAble Devices, Inc. This organization has also developed a system where two users, each with their own PHANToM, can cooperate in a shared virtual workspace, though this system is not yet commercially available.

SensAble Devices is currently engaged in developing additional software support for PHANToM, investigating issues in two-fingered grasping, and developing a system that supports force feedback for a third finger. A low-cost version of PHANToM intended for the mass consumer market is also under development and expected to become commercially available within the next couple of years.

6.2.2.10 SAFiRE

EXOS, Inc. developed and market a sensing and force reflecting exoskeleton (SAFiRE) that applies forces to the thumb, index finger, and wrist. The Phase II SAFiRE device has eight active DOFs: 3 DOFs on the thumb, 3 DOFs on the index finger, and 2 DOFs on the wrist. After investigations to determine a suitable mechanical system, linkages grounded to the forearm that apply 3-D Cartesian forces to the fingertips and palm were chosen. The endpoints of the manipulators for the thumb and index finger are attached to the fingertips, and include passive freedoms that allow for comfortable finger motion.


The SAFiRE device is actuated by DC motors that are remotized and connected to the device joints with a cable and gear transmission. Incremented optical encoders are attached to motors for measuring motor position and an Ascension Bird tracker is attached to the forearm to sense forearm position and orientation. Actuator packages are grounded to the forearm; tips of the manipulator are attached to the fingertips and to the palm via a cuff.

The basic electronics and control hardware consists of two dedicated high performance processors and a Silicon Graphics workstation. One processor handles dynamics simulation software and communications, while the kinematics and low-level device control software resides on the second processor. The VE display graphics and user interface module is implemented with the Sense8 package on the SGI workstation and communicates with the processors via parallel I/O on a SCSI bus. The dynamics simulation environment contains a number of objects that can be manipulated by the SAFiRE device. The kinematics software contains kinematic models of the SAFiRE device and the human hand, and is responsible for transforming joint angles and position information into Cartesian space for the dynamics and graphics modules, as well as converting Cartesian forces to motor torques.

For desktop applications, an optional boom is available to counterbalance the weight of the device. A photograph and some specification details are given in Figure 86. SAFiRE devices are built to order and no general price information is available.

6.2.3 Current R&D

As with tactile interfaces, development of force feedback devices for teleoperation systems has provided the initial starting point for the technology discussed in this section. Indeed, there are many aspects in which the technology used for providing force feedback in VEs is the same as that used for teleoperation.



Specification		
Force/Torque	Thumb	1 lb max
	Index Finger	1 lb max
	Wrist	2 lb max
Joint Motion	Each force reflecting flexion DOF	
	90°, finger radial/ulnar deviation	
	30°, circumduction of thumb	
	45°	
DOF	Thumb	3
	Index Finger	3
	Wrist	2
Friction	Approx. 4 oz-in or less	
Backlash	<1°	
Device Weight	<4 lbs	
Interface	Custom	
Photo courtesy of EXOS, Inc.		

Figure 86. SAFIRE

The remainder of this section discusses the work of individual research groups. While it is unlikely that the identified efforts are the only ones currently underway, they do form a representative set. In addition to those presented below, there were two research efforts for which detailed information was not available. The first was the development of an arm exoskeleton system for the presentation of force feedback by Dr. Bergamasco at the Sculo Superiore in Pisa, Italy. The second is Dr. Robert Anderson's development of software libraries that support the creation of virtual forces. Dr. Anderson is with Sandia National Laboratory.

6.2.3.1 Boeing Computer Services

Led by Dr. William McNeely, researchers at Boeing Computer Services are pursuing the development of a high-fidelity force feedback capability that is scalable to the range of body motion and work volumes encountered in the simulated design, manufacture, and operation of aerospace vehicles. This work employs the concept of robot graphics in which forces are served by robotic mechanisms that are not attached to the body.

In 1994 Boeing researchers conducted proof-of-concept demonstrations of this approach, illustrating how it might be applied to control panel prototyping. As shown in Figure 87, the user, wearing a HMD, stands before an empty physical panel with a rectangular grid of holes. The HMD displays a control panel designed using four different types of controls, and this image is 3-D registered with the physical panel. Whenever the user reaches out to contact a virtual control, the associated hand motion is detected and the contact intention deduced. A robot then quickly moves a physical control of the right type to the anticipated point of contact, pushing it through the hole in the physical panel and holding it there to satisfy user contact. Although only one finger was videometrically tracked, the entire hand received appropriate force feedback, for example, in pushing a button or turning a knob.

Figure 88 shows the first-person view of this process. The layout of controls can be edited by touching controls and special command points. This system demonstrated a natural and effective VE interface and validated the robot graphics approach. Full immersiveness was not achieved with the available apparatus, however, primarily because of inadequate haptic-visual registration.




Figure 87. Robotic Graphics Proof-of-Concept System Overview




Photo courtesy of Boeing Computer Services

Figure 88. Robotic Graphics Proof-of-Concept System

provide forces and torques as dictated by the VE. These attach points would be handed off from robot to robot as required to avoid such problems as mechanical singularities and robot-robot collisions.

Another key development area is the robot graphics "middleware" to support the design-driven selection and placement of static display elements, and rapid prototyping/approximation/deployment of active display elements. Work in this area is also addressing such concerns as human motion prediction, multi-robot choreography, and accident avoidance. It is thought that much of this infrastructure could be developed and safety tested using software simulation in an auxiliary VE testbed.

6.2.3.2 Computer Graphics Systems Development Corporation

Computer Graphics Systems Development (CGSD) Corporation researchers also are investigating the potential of robotic graphics. Following a feasibility study conducted

in 1994, these researchers are developing a prototype virtual aircraft cockpit that includes realistic simulation of the forces and tactile sensations of operating instrument panel controls. Physical fixtures are used for the primary controls (throttle and joystick) and a stereo image of the cockpit, including the user's hand and the out-the-window scenery, is presented in a HMD. The virtual cockpit is intended to provide a highly reconfigurable simulator for design verification and, ultimately, for flight training.

Requirements for the Force and Tactile Feedback System (FTFS) were determined by analysis of cockpit videotapes and lab experiments. The FTFS is a robotic positioning system that tracks the user's hand, anticipates which control is to be actuated, and moves an example of the control into position to be actuated. The positioning mechanism has, on a flat panel, various types of switches and knobs representative of instrument controls, including a number entry keypad. The system provides for positioning of controls in 3 axes and is designed to provide the high positioning speeds needed for realistic operation as a simulator. The user signals his intent to operate a control by reaching for that control: tracking his hand and fingers, the computer performs extrapolations to determine which control is desired. The panel does its final positioning in less than 50 msec, so that the control is stationary before the user touches it. The workspace of the prototype FTFS is roughly 48 inches long, 30 inches high, and 6 inches deep. The controls are positioned to an accuracy of about 0.003 inches. Large motors are needed to move the panel quickly, so the controls are inherently stiff, able to resist over 10 lbs of force without perceptible motion. Safety is a major concern of the design and is the first concern of the development process. Redundant mechanisms are being incorporated to prevent user injury and a rectangular coordinate positioning system is used which cannot intrude into the user's space. The positioning device, when complete, could be coupled to any host simulation. It could also be used with varying visual systems. The preferred interface is dedicated Ethernet.

The FTFS is currently under construction with sponsorship from the US Army Simulation Training and Instrumentation Command (STRICOM). Completion of the prototype is expected by the end of 1996.

6.2.3.3 Hokkaido University, Japan

Led by Dr. Shuichi Ino, researchers at Hokkaido University are developing a system to provide force feedback to an elbow joint. For this, they have developed a metal hydride actuator which uses temperature changes in a metal hydride alloy to control the pressure of hydrogen gas in a bellows system cylinder; the pressure is converted into a propulsive force. The actuator is lightweight (300 g) and compact (a cylinder 20.62 mm in diameter). Using a metal hydride alloy of 6 g, it can generate a power of 20 kgf, and lift a load of up to 10 kg to a height of 50 mm with a velocity of 9 mm/sec, the fall time is roughly equivalent. The actuator is noise-free and produces no sudden impact forces. Experimental trials have demonstrated that the display has similar variable compliance to the human elbow, and that this compliance can be smoothly controlled by a computer.




Figure 89. Elbow Force Feedback Display
(Hokkaido University)

The researchers are investigating approaches for mounting a force display based on the metal hydride actuator on a human arm. Using the mounting shown in Figure 89, two psychophysical experiments were conducted to examine the usefulness of such a display (Shimizu, 1993). In the first experiment, the differential limen of the force sensation was measured at 400 gf and the ability of the display to provide smooth force changes to the elbow was demonstrated. In the second experiment, researchers compared the force sensation level produced by the display with that achieved by placing a real object on the forearm. In this

case, the force sensation level difference between the sensation produced by the force display and that produced from a real object was less than the differential limen and, hence, unnoticeable.

Currently, the researchers are investigating parameters that can be used to provide realistic sensations of weight, resistance, and binding.

6.2.3.4 Massachusetts Institute of Technology, Artificial Intelligence Laboratory

Led by Dr. Ken Salisbury, researchers at MIT, Artificial Intelligence Laboratory, are pursuing two main areas of work for force feedback interface technology. The first area is concerned with further development of PHANToM (see Section 6.2.2.9). Here, much of the work lies in investigating how to move from the single point interaction provided by the current PHANToM device to more general paradigms for multiple finger interaction. The researchers are also looking at replacing PHANToM's thimble interface with passive tools and with 1 DOF tools such as power-driven tweezers. A near-term goal for this work is the development of a system that uses two PHANToMs, each equipped with tweezers, for reaching into a virtual scene, grabbing simulated body tissue, and passing the tissue from one pair of pliers to the other; this work is supported by the Advanced Research Projects Agency (ARPA).

The researchers also are looking at using PHANToM as a platform to support additional sensory modalities. One example of such a modality will be the ability to transmit high frequency vibrations to provide the user with object texture information. The researchers are also talking with CM Research Inc. (see Section 6.1.2.4) about using CM Research technology in providing temperature feedback to PHANToM users.

The Navy Air Warfare Center Training Systems Division (NAWCTSD) is supporting MIT in an initial, practical use application of PHANToM force feedback. The system under development will use a virtual electronic board (comprised of rigid objects) and virtual probe to provide training for electronic technicians. A possible second application will be in bomb disposal training.

The other major area of research is in the development of software technologies. Here one effort is looking at haptic rendering, that is, the process of computing and generating forces. The overall goal is to develop a framework that can represent shapes, bulk properties, and multiple object interactions. Focusing on point interactions, the researchers have developed algorithms for rendering object contact forces, contact persistence, and impedance. They have found that if force information is presented with sufficient bandwidth and resolution, they can produce effects that are perceived as tactile sensations and, exploiting this, have also developed algorithms to render object surface properties such as curvature, texture, and friction. Different techniques for rendering overall object shape are being pursued; vector field techniques have already been implemented and the researchers are currently working with a constraint-based method they call the god object method. Algorithms for rendering non-homogeneous materials are under development. These support efficient implementation by allowing rendering to be limited to a local "window" of surface representation data. (Methods for the haptic scanning of surface property data based on force scanning are also under development.) The current methods used in rendering non-homogeneous materials apply the B-spline surfaces geometric modeling technique, though the use of potential field methods is being explored. In the course of this effort, the researchers have developed several demonstration applications, including a virtual control panel, a needle biopsy simulator, a tissue palpation system, and a simple virtual world where two PHANToMs are used to manipulate building blocks. Ultimately, this work will result in a haptic renderer that can accept CAD data as input and so facilitate construction of virtual worlds that exhibit a range of object properties and object interactions. The current prototype renderer supports haptic rendering for simple stationary objects. Future work on object rendering is expected to look at potential energy function representations.

Another effort concerned with software technologies involves the development of a reduced fidelity model for compliant objects that can be run in real time. Finally, as part of the ARPA task, the researchers are defining models of human organs and tissues that can be used to provide information about how these body parts "feel." Once the model parameters and structures have been determined, data will be collected to build a library of models that can be customized to particular patients as needed.

6.2.3.5 Massachusetts Institute of Technology, Department of Mechanical Engineering

At MIT, Department of Mechanical Engineering, researchers under Dr. Ian Hunter have been undertaking a large-scale effort to develop a system for eye surgery that exploits the capabilities of both teleoperation and VEs¹. The system is being developed as an exper-

imental testbed that could be used to study the effects of feedforward and feedback delays on remote surgery. It also is intended for use in research on how mechanical and visual telepresence can enhance the accuracy and dexterity of microsurgeons.

The teleoperation part of the system allows a surgeon to perform surgery using a teleoperated microsurgical robot (MSR-1) master and slave. Visual, mechanical, and auditory information is exchanged between the master and slave. With respect to visual information, the surgeon uses an HMD to orient the stereo camera system observing the surgery and to feedback images from the camera system to either the HMD or an adjacent screen. The surgery is performed using pseudo tools supported on active limbs mounted on a mechanical master. The surgeon's movement of the pseudo tools is scaled down, by 1 to 100 times, and transmitted to the microsurgery tools operated by the mechanical slave. Forces acting on the limbs of the MSR-1 slave are transmitted back to the surgeon, scaled up appropriately, via the pseudo tools on the MSR-1 master. Additional information is provided audially as a stereo tone whose amplitude and/or frequency is a function of the forces exerted at the tool-tissue interface. The computer system that controls the equipment enhances and augments images, filters hand tremors, performs coordinate transformations, and performs safety checks. The primary computer hardware is two IBM RISC System/6000 workstations.

The mechanical parts of the MSR-1 system, master and slave, employ six direct-drive rotary electromagnetic actuators for each active limb. Each set of six actuators is arranged in a redundant parallel configuration that supports motion in three linear and two rotary DOFs. Force and displacement transducers are integrated into the actuators. Figure 90 provides some details about the mechanical master and slave, including a photograph of the master device. The VE part of the system centers around a detailed continuum model of the eye anatomy, mechanics, and optical properties, supported with a less detailed geometric/mechanical model of the face. Mechanical finite element models of structures in the eye make it possible to calculate and display the deformation of tissue as it is manipulated, and to calculate forces to be feedback to the mechanical master. Together with a computer simulation of the MSR-1 slave, these models provide a training environment where the surgeon can practice with the use of the MSR-1 system and plan surgical procedures. Active mannequin faces are used for testing the microsurgical system and training surgeons in its use. The VE also plays a role here in providing input to the machining process used to construct a mannequin face.

Currently, further work on MSR-1 is awaiting the additional investment needed to support refining the existing prototype system and conduct testing in a real surgical environment. Meanwhile, these researchers, funded by ARPA, are developing a system to support heart surgery, HSR-1. The new system is closely based on the MSR-1 and will have

¹ This work was started at McGill University, Department of Biomedical Engineering, in cooperation with the School of Physical and Occupational Therapy at McGill University, and with the departments of Mechanical Engineering and Engineering Science at the University of Auckland, New Zealand.




Figure 90. MSR-1 Mechanical Master/Slave

similar support for force feedback to the surgeon. It is currently in the early stages of development.

6.2.3.6 McGill University, Canada

Dr. Vincent Hayward at McGill University, Research Center for Intelligent Machines, is leading researchers in the development of haptic devices. One of these is a 6 DOF, optionally 7 DOF, force-feedback interface device intended for use in VEs and teleoperation. Called the Stylus, this device derives its name from its intended use as a small handle held with a precision grasp. In order to determine the requirements for such a device, these researchers have undertaken a number of experiments designed to provide insight into how well the human hand can discriminate the direction and nature of small motions (Hayward, 1995). Using a specially designed 6 DOF stimulator able to vary the amplitude, direction, and nature of small motions, subject sensitivity was found to vary greatly with training. In certain conditions, some subjects were capable of discriminating between two consecutive motions vibrating across orthogonal directions at up to 80 Hz, frequencies after which the motion was perceived as simply vibration.

Studies with another specially designed piece of equipment, a 2 DOF haptic device named the Pantograph, were conducted to determine an appropriate workspace for the Stylus. Results here led the researchers to decide on a work volume of 10 x 10 x 10 cm and an angular workspace on the order of 90° of pitch and yaw, with a roll of 180°. Additional studies with the Pantograph intended to characterize device factors that best transduce sensations of shock and hard contacts, found that the acceleration capability was probably the most useful factor. Additional requirements for the device were derived from the literature. These include requirements for a wide frequency response, high levels of accuracy in presenting forces, mechanical impedance that can be programmed over several orders of magnitude, and precise dynamic response up to 50 - 100 Hz. For its physical structure, the

desktop Stylus device uses a single stage design with grounded actuation coupled by a combination of polymeric tendon transmissions and linkages to the active end. A separate actuator pack employs conventional electric motors. Custom-designed sensors based on optical techniques are used to measure displacement and forces.

A prototype device, see Figure 91, has been developed. This device provides 3 DOFs for displacement, 3 DOFs for handle orientations, and 1 DOF for pinching motions. Current work is involved in developing a driver for the Stylus that supports an analog interface to a PC. Other ongoing work involves further investigation of human factors and ergonomic issues for haptic interfaces. This work will include experiments investigating, for example, specific frequency and force requirements for particular applications.



Figure 91. 7 DOF Stylus (McGill University)

Currently, McGill University is working with MPB Technologies Inc. in commercializing the Stylus under the name Miniature Hand Controller. MPB Technologies, Inc. are also preparing a commercial product of a highly similar device, called the Wide-Span Hand Controller, that will provide a 30 cm diameter sphere workspace.




Figure 92. Pantograph (McGill University)

The Pantograph, mentioned previously, was developed in early 1993 as a collaborative effort with Christophe Ramstein from the Center for Information Technology Innovation, a research organization that is part of Industry Canada located in Laval, near Montreal. The Pantograph originally was designed as a computer interface in an effort to provide visually handicapped people access to graphical applications. A number of prototypes are operating at various laboratories in Canada in projects ranging from a surgical aid to human factors research. In one case, it is currently being compared with a conventional input device for use in microgravity environments. A picture of the Pantograph is shown in Figure 92.

In other work, Dr. Hayward, in collaboration with Dr. Raymond Hui, is supporting the Canadian Space Agency in the development of haptic interfaces for use in space and terrestrial applications. The design of these 3 DOF devices based on parallel linkages is

motivated by the search for high structural transparency and simplified kinematics. Current work is looking at combining several of these low DOF devices, one per finger, to avoid the inherent complexity of higher DOF devices. Dr. Hayward is also proposing the development of a standard specification for haptic interfaces.

6.2.3.7 Ministry of International Trade and Industry, Agency of Industrial Science and Technology (MITI/AIST), Japan

Researchers at the MITI/AIST, National Institute of Bioscience and Human Technology, are developing a CAD system where the shape of 3-D surfaces can be evaluated using surface tracing and localized lighting schemes. Led by Dr. Yukio Fukui, this work started with modifying a conventional XY-recorder into a 2-D force feedback device. This device provided a 20 x 20 cm workspace in the horizontal plane with a 4 Hz response.

Experiments looking at the effectiveness of this device for virtual shape recognition have been conducted (Fukui and Shimojo, 1994). Four subjects were asked to perform discrimination tasks using either visual only (presented on a computer screen), haptic only, or both visual and haptic displays. The first task required selecting true circles from a series of deformed circles, where all the circles were visually presented as superimposed on a series of concentric squares. The second task required selecting pairs of lines that formed a straight line from a series of pairs that were jointed at angles other than 180°, here the lines were visually presented as superimposed on a background of lines connected at an angle. The results of the experiments showed that optical illusions commonly occurred when visual feedback was present, and the haptic only feedback condition gave a better mean performance than visual only feedback. When both types of feedback were provided, the mismatch of the visual and haptic information led to a performance similar to that achieved for visual feedback alone.

More recently, the researchers have developed a Cartesian 6 DOF force feedback manipulator. To achieve the desired stiffness, toughness, linearity, and economy, the manipulator uses bowl screws on a parallel movement for the x , y , and z dimensions, and worm wheels for rotation movement about these axes. The manipulator is driven using an AC servo motor to provide high speed and precise positioning (0.01 mm/pulse) and an external potentiometer is used to measure manipulator displacement. The forces exerted by the user are measured by a force sensor on the end-effector. The workspace provided is 30 x 30 x 15 cm with 200° about each rotational axis. Figure 93 shows a photograph of this device.

With the addition of a programmable mechanical impedance control, the 6 DOF manipulator is being used in studies investigating the spatial manipulation capabilities of the human hand. Additionally, it is being used as a 3-D input device for a CAD environment and early experiments have demonstrated its value for deforming virtual surfaces. Current work also is focusing on incorporating adaptive damping in the feedback loop to control device vibration and unstable movements.

6.2.3.8 Northwestern University


Led by Dr. Ed Colgate, researchers at Northwestern University, Department of Mechanical Engineering, are looking at several aspects of force feedback for interaction with VEs. They have developed a 4 DOF force feedback device and are currently investigating the issue of stability with respect to force feedback interfaces.

The 4 DOF force reflecting manipulandum provides a user with a joystick-type interface that can be mapped to a virtual hand tool in a VE, allowing the user to perform various mechanical tasks. Developed as a research tool, the device is designed to generate high impedances so that it can support the representation of such things as stiff springs and hard walls in the VE, as well as to exhibit low inertia and low friction. It uses direct-drive, parallel structures with closed-chain kinematics to provide translation in three directions and rotation in the horizontal plane. The workspace is free of singularities. High resolution optical encoders on the brushless DC-motor actuators and precision potentiometers on the non-actuated joints sense the angular position of the device joints to provide endpoint position and orientation of the endpoint (handle). Safety concerns are handled by mechanical stops and covers, hard-wired enable/disable circuitry and accelerometers, and software control. A photograph and some specification details on the device are given in Figure 94.



Photo courtesy of MITI/AIST, National Institute of Bioscience and Human Technology

Figure 93. Cartesian 6 DOF Force Feedback Manipulator (MITI/AIST)



Specification	
Peak Force	20 lbs
Continuous Force	10 lbs
Output Torque (Continuous)	12 in/lbs
Translational Range of Motion	Circle with 7 in diameter
Rotational Range of Motion	$\pm 45^\circ$
Vertical Range of Motion	3.5 in (not used)
Update Rate	250 - 1000 Hz
Latency	3 - 4 msec

Figure 94. 4 DOF Force Reflecting Manipulandum (Northwestern University)

The device is supported on a Inmos T805 Transputer with a VME bus backbone. Accordingly, the computations necessary to support the VE simulation and its user interface are performed on a network of distributed memory parallel processors. Four of these processors are dedicated to control of the force feedback device, performing the calculations necessary to determine endpoint position and velocity, and to determine the motor torques necessary to produce feedback endpoint forces and torque. The manipulandum is currently on loan to the National Aeronautics and Space Administration (NASA) Johnson Space Center, where it is being used to examine the feasibility of using VEs in the training of extra-vehicular activity (EVA) procedures.

The major thrust of the researchers' current work is in developing a physics-based approach for haptic interfaces that ensures that a VE simulation is governed by the conservation laws that operate in the real world. Using a 1 DOF force feedback device developed for this purpose, they are focusing on the requirements needed to guarantee stability when a user interacts with a VE. Although no method for guaranteeing system stability has yet been found, their work has shown the need for inherent physical damping in the haptic interface to increase its passivity, and digital filtering of the velocity signal to achieve high values of virtual damping. They have also found that high update rates increase the achievable stiffness of virtual walls. (Based on these findings, dampers have been introduced to the 4 DOF manipulandum already discussed.) Examples of current work in the area of stability include the definition of non-conservative stability conditions for systems involving unilateral constraints, and the development of methods for real-time simulation of multi-body systems guaranteeing physical passivity. As part of their efforts in this area, the researchers are developing a haptic programming language that will facilitate the development and modification of VEs that provide physics-based haptic interactions.

As a first step in a study aimed at developing a theory of tool use to guide haptic interface, VE, and telemanipulator design, the researchers have conducted an experiment investigating the impact of environment damping (Millman and Colgate, 1994). Subjects were asked to position and maintain contact with a target region in a VE, where the target was distinguished by a rise or drop in the ambient virtual damping. The damping was simulated by the impedance of a 1 DOF manipulandum. Three visual feedback conditions were used: visual feedback was provided by showing the position of the manipulandum handle as a cursor on a screen, no visual feedback, and visual feedback showing the position of both the handle and target region on the screen. For very large differences in target and ambient damping, the results showed that haptic feedback alone gave performance almost equivalent to that attained when subjects could see the position of the hand and the target region on the screen. The subjects could detect the target region when the differences in the ambient and target virtual damping was greater than 2.27 N/m, though the researchers expect a greater level of discrimination is possible with equipment characterized by lower system noise.

In future work, these researchers hope to put the experience they have gained with force feedback interfaces to use in the development of a 6 DOF device.

6.2.3.9 Rutgers University

For several years, Dr. Greg Burdea has led a group of researchers at Rutgers University, Center for Computer Aids for Industrial Productivity, in the development of force feedback interfaces for grasping and manipulating virtual objects. The work has focused on the development of a portable dextrous hand master that has evolved through the Portable Dextrous Master with Force Feedback (PDMFF), the Rutgers Portable Force Feedback Master (RM-I), to the current Second Generation Rutgers Master (RM-II).

The central element of the RM-II design is the placement of four custom-designed pneumatic micro-cylinders on an "L"-shaped platform positioned in the user's palm, the whole being mounted on a thin leather glove. The actuators are additionally attached to the tips of three fingers and thumb, as shown in Figure 95, to deliver forces to those points. Attachments are by velcro strips to accommodate various user hand size. Position sensing for the fingers is integrated into the device by means of an assembly of two Hall-effect sensors mounted on the platform, an additional Hall-effect sensor mounted on each fingertip, and an infrared LED-phototransistor pair placed within each cylinder. A Fastrak position sensor mounted on the back of the hand provides wrist position and orientation. A separate interface box is used to house the proportional analog servo-controllers used to regulate the air pressure and, hence, the forces applied to each fingertip.




Figure 95. Second Generation Rutgers Master

Software supporting the RM-II has been developed to integrate this device (or the RM-I) into a complete VE interface. The VE system that currently uses this interface is hosted on a set of workstations connected via Ethernet and supports StereoGraphics LCD glasses. The VE itself consists of a room with perspective grids, a virtual hand, and a selection of virtual objects (ball, soda can, and spring). Object stiffness can be adjusted to model both elastic and plastic objects. Forces are calculated using Hook's Law, and gravity is modeled to allow objects to bounce off the walls. Grasped objects deform graphically when squeezed by the virtual hand. Present research is aimed at accommodating multiple objects

in the same scene. The researchers plan to port this software to the Sense8 WorldToolKit Version 2.0 and to develop a stand-alone PC architecture for the force feedback interface.

Using this VE system, and the earlier RM-I, the researchers conducted a series of human factors experiments to test the usefulness of force feedback for VEs (Gomez, Burdea, and Langrana, 1995). In one experiment, the task was to grasp and manipulate a deformable ball without indenting it more than 10% of its volume. Eighty-four subjects participated, divided into six groups. Each group was provided with a graphical representation of the ball deformation plus some combination of force feedback, visual bar-chart of output pressures, and auditory displays. (The auditory feedback, presented through headphones, provided a sound frequency proportional to the current deformation of the ball and was used as a substitute for tactile feedback.) The first half of each group performed the task using a monoscopic display for the graphics feedback, the second half used a stereo display and active LCD glasses. Among the non-redundant feedback modalities, force feedback produced the best result. When redundancy was present, the force and auditory feedback combination was superior. Present experiments are aimed at repeating the trials using the newer RM-II, in order to compare it with RM-I performance.




Figure 96. Virtual Knee Palpation System

VE itself supports three types of objects: a virtual hand, virtual knee, and room walls; the user's viewing perspective of the rendered scene can be adjusted using a 6-D trackball. A photograph of this VE is shown in Figure 96. Recent work has focused on the development of the collision detection and deformation algorithms and the proof-of-concept system currently supports one finger palpation of tissue compliance. The computational load imposed by shape modeling and calculation of force data has resulted in a low frame rate, between 3-5 frames/sec, and ongoing work is improving the simulation to allow increasing the frame rate to a minimum of 14 frames/sec, and providing a force bandwidth of 10-12 Hz.

Additional, and future, research is pursuing complex multi-fingered manipulation and the development of an RM-II application to support explosive handling in suitcase inspection at transportation facilities.

Work is also underway to develop more general software support for the RM-II. This software is currently being used to support force feedback in a system for virtual knee palpation that can be used by surgeons to train and plan for knee surgery. For this application, geometry data for a complete knee joint was modified to support tissue deformation. A DataGlove is used for position sensing and the RM-II for tissue manipulation. A Sun 4/380 is dedicated to handling the force feedback for the RM-II, with the rest of the simulation being performed on a HP 755 CRX workstation. The

6.2.3.10 Suzuki Motor Corporation

At Suzuki Motor Corporation, Technical Research Center, researchers led by Dr. Yoshitaka Adachi are investigating the use of force feedback for manipulating the free-form surface of virtual objects. They have developed a simple distribution function algorithm for recognizing the interference between a 3-D cursor and virtual objects, and calculating the direction of reaction forces in real time. These forces are then generated using impedance control. The force feedback device they have developed to display the forces is called SPICE.

SPICE is an articulated mechanical structure with invariant and decoupled arm inertia and 6 DOFs. Each joint is driven by a direct drive DC motor capable of providing a wide range of torques. The arm structure and its mass distribution have been optimized through an evaluation of arm dynamics with a generalized inertia ellipsoid. Optical encoders are used to sense joint angles and so provide information on the position and orientation of the user-held grip. The grip provides the user with the means to interact in a VE, and a 3-D cursor is used to represent the user's hand. One microcomputer with a floating point co-processor is used to simulate the virtual haptic environment. Another microcomputer, with floating point co-processor and three vector processors, controls the force feedback device. The device and the various computers communicate via a VMEbus. A photograph and specification details for SPICE are given in Figure 97.




Figure 97. SPICE

SPICE has been used in evaluating an approach for representing free-form stiff virtual objects. One of the limiting factors for the generation of stiff surfaces is sampling rate and this approach focuses on reducing the computational requirements needed for detecting collisions with virtual objects, so that the sampling rate is not reduced to unacceptable levels. Based on the position of the user's fingertip in the VE, a tangential plane that includes the point on the virtual object in the VE that is nearest to the fingertip is defined. The position and orientation of the virtual plane is updated at low frequency to accommodate finger movement. At the same time, collision between the fingertip and virtual plane is checked

and, if necessary, a reaction force is calculated. The primary advantages of this approach are that it simplifies calculation of force vectors and allows detection of collisions to be computed independently from the impedance control of the force feedback device. Using SPICE, the researchers conducted an experiment to determine the frequency of impedance control required for effective generation of force sensations; for tracking on a stiff virtual wall, Adachi, Kumano, and Ogino (1995) report that an impedance control greater than 500 Hz was needed. The researchers also investigated the update rate required for the virtual plane when representing a curved surface in the VE. This experiment employed a virtual cylinder with a 75 mm diameter, an impedance control of 1,000 Hz, surface stiffness of 10,000 N/m, surface viscosity of 1000 N/m/sec, and artificial friction of 600 N/m/sec. A virtual plane update rate of 3.3 Hz gave the impression of a smooth curved surface when finger velocity was around 20 N/m/sec. A lower update rate (2.5 Hz) gave a feeling of a bumpy surface unless finger movement was slowed to 8 mm/sec.

A separate set of experiments, using SPICE, has been conducted using sensory evaluation methods to determine the impedance characteristics that make virtual push-buttons comfortable to operate. More specifically, these other experiments have examined the effect of physical parameters such as spring stiffness and damper viscosity for button pushing. The virtual button was designed as a massless plate backed by one spring, with a second spring and damper positioned under the first spring to provide the feeling of the bottom of the button. The first experiment in this series examined the feeling of the bottom of the button in terms of the sensory factors "stiffness" and "evaluation." Here Adachi reports that in the case of low damper viscosity, the stiffness factor increased linearly with increases in viscosity and with increases in spring stiffness. At higher levels of damper viscosity, the stiffness factor was influenced by the viscosity rather than spring stiffness. No clear relationships between the evaluation factor and spring stiffness or damper viscosity were found. The second experiment examined the operational feeling of push-buttons, with respect to initial load and spring stiffness. It was found that the stiffness factor score increased linearly with increases in the initial load, buttons with a small initial load lacked of feeling of crispness in the button surface, and users did not like buttons that needed a large force to operate them. The final experiment in this series assessed the effect of incorporating a "click" feeling into the buttons. This sensation was introduced by increasing the reaction force as a button was pushed and at a given maximum value, in the middle of the push stroke, quickly decreasing the reaction force. The experimental results demonstrated that this type of clicking decreased the required operational force while keeping a crisp feel at the surface of the button.

Currently, SPICE is being used in a project to develop advanced user interfaces for CAD and 3-D modeling systems. One of the applications under development in this project is a 3-D sketchbook intended for use by plastic surgeons in preparing for orthopedic surgery.

6.2.3.11 Tokyo Institute of Technology, Japan

Researchers led by Dr. Makoto Sato at the Tokyo Institute of Technology, Precision and Intelligence Laboratory, have developed a force feedback device called the 3-D Spatial Interface Device for Artificial Reality (SPIDAR). For this device, the user inserts his index finger into a cap that is held by four strings. Using a system of pulleys and motors, string lengths provide a means for measuring finger position, and the tensed strings provide for the presentation of force sensations at the cap. SPIDAR is controlled by a DX4-100MHz PC. A photograph and some specification details are provided in Figure 98.




Figure 98. SPIDAR

The utility of the force feedback provided by the SPIDAR was examined in an experiment where three subjects used SPIDAR to deform a virtual cup-shaped object into an object shaped like a soccer ball (Ishii, 1994a). In the absence of force feedback, the subjects found it difficult to complete the task and often gave up; with force feedback, the task was completed within a few minutes.

A version of SPIDAR, SPIDAR II, that provides force feedback to the thumb and index finger has also been developed. Two initial experiments that looked at the effectiveness of SPIDAR II for pick-and-place tasks have been performed (Ishii, 1994b). For these experiments, the VE was generated by a Silicon Graphics INDIGO2 XZ, with a visual display presented on a screen and viewed by a user wearing stereoscopic glasses, and SPIDAR II providing force feedback. In this configuration, the refresh rate for the force generation was 100 Hz and forces ranged from 0 - 4 N, with an incremental step of 0.016 N. The experimental task was to position a 5 cm³ block on a target circle marked on a raised platform. For the first experiment the block weighed 50 g, and for the second experiment the weight was varied (20, 35, 50, 70, 100, and 150 g). Three subjects were used for each experiment. The results of the first experiment showed that the provision of force feedback did not greatly impact overall task completion time. The pick-up and positioning parts of the task were performed more quickly with the guidance provided by force feedback, but the moving part of the task was slower with force feedback than without it, perhaps due to the effort required

to move a virtual object that possessed weight. The force feedback did tend to result in more accurate positioning of the block on the target (accuracy was measured as the distance between the center of the bottom of the block and the top of the platform). The results of the second experiment showed that increasing the block weight slowed task completion time; as in the real world, a heavier object is more difficult to move and position than a lighter object. The optimum block weight was between 35 and 50 g. Additional experiments are being conducted to study the representation of object weight in a VE.

With respect to providing force feedback to both hands, ergonomic experiments designed to discover appropriate positioning between a user and SPIDAR devices have been performed. Based on the results of these experiments, a new version of SPIDAR is being developed that allows the use of both hands in combining 3-D objects to construct new objects. Initial usage has shown that this system does, indeed, support users in tasks that require two-handed manipulation directly, or the use of two hands in cooperation. Current work includes enhancing this system to support delicate operations and providing auditory feedback.

A closely related effort is focused on the use of SPIDAR to support a virtual workspace for the collaborative design of 3-D objects. Accordingly, this virtual workspace supports both face-to-face interaction between a pair of participants and interaction between a participant and an object. It is structured around what the researchers term a "dialog space" and an "object space," with participants switching their attention between these spaces as required. The specific design requirements imposed for this system are: (1) direct manipulation using both hands, (2) provision of force feedback, (3) support for pick-and-place operations on objects, (4) a wide range of hand motions with a significant number of DOFs, and (5) easy and safe operation, using an inexpensive system. A prototype system with single hand direct manipulation has been developed, supported by a local area network. Each participant is provided with a SPIDAR and two screens in juxtaposition: one screen for the face-to-face communication between the participants, and the second for display of the VE where 3-D objects are designed. While only one participant can actually manipulate an object at any one time, both can be in contact with the object so that the second participant can feel the forces exerted by his partner. Microphones and audio speakers are used both to support voice communications and to present sounds of object collisions. In a small-scale experiment, four pairs of participants performed a hand-over task, where a virtual block is passed from by one participant to his partner. The force feedback provided by SPIDAR enabled quick and accurate passing of the block. There were no instances of the block being dropped, as frequently occurred when no haptic feedback was provided. More recent and ongoing work in this area is focusing on multi-user collaboration, support for multiple virtual objects, and networked force feedback interactions in the presence of time delay. A demonstration for networked VE is under development, this is a virtual tennis game using two enlarged SPIDARs (with a working space of 3 m³) connected to each other.

6.2.3.12 University of North Carolina

Some of the earliest work in force feedback displays for VEs was conducted at the University of North Carolina (UNC) at Chapel Hill, Computer Science Department, and these researchers, led by Dr. Fred Brooks, continue to be active in this field. The overall objective of the work is to investigate and develop methods for providing high quality force feedback in real applications.

The Force-Feedback Project, which began in 1967, first focused on the development of a system to support scientific visualization in the area of molecular docking, the Docker application. This application provides graphic (wire-frame) representations of molecules and their inter-atomic forces to allow a user to adjust the relative position and orientation of molecules while searching for minimum energy binding sites. A series of systems have been developed, evolving from a 2-D system, through a 3-D system and a 6-D system for a simple docking task, to a full 6-D molecular docking system called GROPE-III. These later systems have employed a modified Model E-3 Argonne Remote Manipulator (ARM) for force feedback display. (The ARM is a 6 DOF device developed at Argonne National Laboratories for teleoperation applications. It uses a hand-grip display, with joint action at the shoulder and outward. It provides a workspace of approximately 1 m³. Forces are generated by AC electric servo-motors and joint positions are measured using analog potentiometers.) The researchers have made several modifications to the ARM, including the addition of dials for controlling the twistable bonds found in some drugs. In the GROPE-III application, the force-feedback device runs on a dedicated PC with a force update rate of 15 Hz. The visual feedback is generated by a Silicon Graphics Onyx with a Reality Engine, and the displays are presented via a StereoGraphic Crystal Eyes unit. GROPE-III runs in synchronous mode, pausing the simulation as necessary to wait for position measurements to be collected, and the resultant forces calculated and sent to the ARM for display to the user. Figure 99 shows a user working with the molecular docking application.




Figure 99. Molecular Docking Virtual Interface

times when the subjects were just thinking, that is, not manipulating anything, are subtract-

Researchers used GROPE-III for an experiment that looked at the effectiveness of force feedback display in a complex molecular docking task (Brooks, 1990). The subjects for the experiment were twelve experienced biochemists. The results showed that the 6-D rigid-body docking part of the task was about 30% faster with the force feedback as opposed to only visual feedback, and drug trajectory paths were 41% shorter with force feedback. However, while the overall elapsed time performance with the force feedback was improved, the difference was not significant. The reason for this is believed to be due to the large amounts of thinking time the task required. If the

ed out, the overall time for the 6-D docking task was 1.75 times faster with force feedback than without. This, and other, experiments have shown that force feedback can facilitate the performance of molecular docking tasks. However, the researchers believe that the major contribution of GROPE-III lies in its ability to give biochemists deeper and new insights into molecular docking issues.

Recently, the researchers have installed a custom-designed PHANToM (6 DOF position sensing, 4 DOF force feedback) as an alternative interface to the molecular docking system. The system now is primarily used at the UNC for demonstration purposes, and is in occasional use by UNC biochemists. It has also been installed at Wright Patterson Air Force Base (WPAFB), Materials Laboratory, where it is being used to investigate the packing of molecules in liquid crystals. (The WPAFB implementation uses a Cybernet PER-Force arm for force feedback.)

Another application under development at UNC supports manipulation of flexible molecules. The initial version of this system, called SCULPT, does not provide force feedback. It is marketed by Interactive Simulations Inc. and in use at Duke University to support the design of amino acids that can hold molecules to required shapes. UNC researchers are extending the system to include force feedback using the ARM and PHANToM.

One of the current focuses of the Force-Feedback Project has been the development of a software library to support the use of force feedback devices. This library accepts force inputs specified in Newtons and torques specified in Newton-meters, and provides position information in units of meters in Cartesian coordinates and orientation in radians. It is designed to be device-independent and the Application Programmer Interface is intended to support its use by application developers unfamiliar with force feedback technology. The library has a client-server structure so that the client portion can run on the same computer as the application and the server portion run on a computer dedicated to supporting the force feedback device. Currently both the ARM and PHANToM devices are supported.

A second major area of research has been the Nanomanipulator (nM) application, being performed as a collaboration between UNC Departments of Computer Science and Physics, and the Department of Chemistry at the University of California, Los Angeles. The goal of this work is the fabrication of nanometer-scale structures in the study of materials relating to quantum effect devices. Here the ARM, and more recently PHANToM, have been used in supporting a range of scanning probe microscopes. The initial implementation provided a VE interface to a scanning tunneling microscope, and later implementations have supported an atomic force microscope. As a real-time visualization application, these systems present a rendered 3-D color surface image to the user, who can then view and feel the surface representation using the force feedback interface and, with the atomic force microscope, make direct surface modifications. Operation of the atomic force microscope virtual interface is illustrated in Figure 100. The display tools support surface scaling and grabbing, flying, and lighting adjustments. Measurement functions are supported by allowing the user to superimpose a reference grid on the surface. Standard VCR-like controls

allow the capture of data for off-line analysis. The suite of surface modification tools support area or selective sweeps, line drawing, and engraving.



Figure 100. Operation of the Atomic Force Microscope Virtual Interface

been used for the dissection and movement of the Tobacco Mosaic Virus. Current work with the system is focusing on better surface representations using the PHANTOM device.

The effectiveness of the nM VE interface with force feedback has been demonstrated in several instances. In its first test, researchers investigating a graphite surface were able to identify sheets of graphite tilted up out of the surface that had been unnoticed using conventional visualizations. Subsequently, use of the nM led to the discovery of a new mechanism for surface modification. In an experiment in the manipulation of colloidal gold particles, users were able to maneuver selected particles without disturbing the surrounding material in a short time, actions which may not be possible using conventional tools. The nM also has

6.2.3.13 University of Tsukuba, Japan

Dr. Iwata at the University of Tsukuba, Institute of Engineering Mechanics, developed the HapticMaster, a desktop force display now commercially available from Nissho Electronics Co., see Section 6.2.2.5. Researchers at the university are using this device to develop a force feedback environment that supports the design of 3-D shapes.

A pen-based device has also been developed and is being used in interactive deformation of free-form surfaces. For this application, the user is provided with a toggle switch that allows the pen to be put in a special mode that adjusts the position and orientation of the surface, and a slider that is used to select the appropriate deformation area. The reaction forces are displayed vertically to the original surface and increase proportionally to the displacement of the pen point. The researchers developed a sine curved-based deformation algorithm for free-form surfaces and, when using the pen interface to push or pull on the surface while holding down the pen button, the user is able to feel the reaction forces. Use of the HapticMaster for interactive deformation of free-form surfaces also has been demonstrated.

Another example application is volume visualization. Here, visual and force feedback are integrated to support multi-dimensional representation of volumes. This system includes both a HMD, with graphics generated by a Silicon Graphics IRIS INDIGO2, and the force display. The pen-based device acts as a 3-D pointer to provide user input to the system. In an experiment to evaluate the effectiveness of the force feedback interface in a volume classification task, three subjects were asked to count the number of high density

cores (1 to 5) embedded in a less dense volume, and then point out each such core. The results showed that the provision of force feedback doubled the accuracy of the pointing task. A second experiment using 6 cores and 4 subjects gave similar results. Current work in this area is examining different methods for mapping voxel data to force and torque, cross influences between force and torque sensations, and further applications of volume haptization.

Dr. Iwata is also engaged on the development of a VE interface that will permit a user to walk in virtual space; this work is discussed in Section 7.2.2.7.

6.2.3.14 University of Washington

The University of Washington, Department of Electrical Engineering, has three separate efforts underway that are concerned with providing force feedback for VEs. One effort is developing a pen-based force display, another is developing a high bandwidth force display, and the third effort (being performed in conjunction with NASA, Johnson Space Center,) is investigating the use of robots in providing force feedback. This work is under the leadership of Dr. Blake Hannaford.

The prototype pen-based force feedback device has been developed as a tool for precise manipulation in a VE, or for scaled telemanipulation. The design goals for the device were driven by the requirements for medical microsurgery. Chief among these were requirements for no backlash or lost motion, minimum friction, and minimum inertia. The user can use the device via fingertip or with a pointed object such as a pen or scalpel. The actual device is a 3 DOF direct-drive parallel manipulator structured as a 2 DOF actuator-redundant parallel cartesian robot with the third DOF by provided by a rotary joint for vertical movement. It is driven by flat coil actuators taken from the read-write heads in hard disk drivers. A photograph and further details are provided in Figure 101. The device is sup-




Figure 101. Pen-Based Force Display (University of Washington)

ported by a real-time controller for PC 486 machines and the resulting force feedback sys-

tem has been used in a VE testbed to demonstrate the use of force feedback in touching virtual objects. A portable library of basic polygonal objects has been developed, together with models for object interaction. Future goals for this work include the development of a VE composed of 3-D solid objects that can be sensed and manipulated using the force display, and characterizing the mechanical impedance of the human hand, when using the device, as a source of mechanical noise. The researchers are also continuing development of the 2-D component of the device.

The high bandwidth force display is a 2 DOF device intended for use in studying simulated interaction with heavy and stiff virtual objects involving whole-arm motions. Kinematically, this device consists of a simple cartesian mechanism driven by brushless DC motors through steel cable transmission. The user holds a knob mounted on the device end-point which is decoupled in orientation. A

Specification	
Peak Force	Up to 1200 N
Continuous Force	400 N
Position Resolution	0.015 mm
Sampling Rate	1 kHz
Static Friction	3.5 - 4.0 N
Workspace	3.04 x 4.06 m

**Figure 102. High Bandwidth Force Display
(University of Washington)**

specification for the device is given in Figure 102. This device is supported by a real-time controller for PC 486 machines. Current software can simulate arbitrary environments composed of polygons and circles. Future work with the high bandwidth force display is expected to include the addition of another motion axis and the development of a more compact version of the device with lower static friction.

Dr. Hannaford is just starting a new project where University of Washington researchers will work with NASA, Johnson Space Center, scientists in developing a robotic graphics system to support EVA training for astronauts. This work will extend an existing VE already in use for EVA training with force feedback via a 7 DOF robot.

6.3 Summary and Expectations

Tactile and force feedback provide important sensory modalities that are prerequisites for many types of practical VE applications. Without these modalities, applications that require complex or precise interactions with the environment, or between users who are not physically present in the same location, are not possible. As yet, however, haptic interface technology can support only very limited types of tactile and force feedback.

Several groups of researchers are investigating the development and use of tactile feedback. Between them, they are addressing the ability of tactile displays to present contact force, slip, texture, vibration, and thermal sensations. Several prototype devices have been developed for experimental purposes and shown that these sensations, as least in primitive form, can be generated. For example, in tests using a blunt pin tactile displays, users were able to discriminate between simple patterns such as a point, line, and plane. Even so, research into tactile feedback is in its initial stages. By and large, current tactile displays provide sensations to a very limited area, usually the fingertip, and many of the approaches

in use will not scale up to provide varying sensations over a larger area. A wide variety of actuator types are being employed, including shape memory alloy, pneumatic, electromagnetic, and piezo-electric technologies and no single technology appears capable of supporting all types of tactile feedback. While each actuator technology has its own particular limitations, they all suffer from relatively large physical dimensions that also constrain their practical use. Studies that identify the most pertinent types of tactile feedback for specific types of applications, and the most appropriate technology for displaying that feedback, are needed. One important practical issue must be to identify where and how trade-offs can be made between the tactile and force feedback modalities. Further evidence of the immaturity of tactile feedback technology is given by the absence of general software models that can be used to determine the sensations that need to be generated with respect to a particular interaction with the environment; with the exception of contact forces, work on developing such models has yet to start. Since tactile sensations depend on a range of physical properties (such as microscopic geometry, coefficient of friction, kinetic elasticity, and thermal conductivity), empirical studies will be important in determining the accuracy that needs to be modeled for practical representation of, for example, surface texture.

Much of the basic psychophysical information needed to support a tactile interface in VEs is available, although there are gaps that need to be filled. There is a lack of data, for example, on the human capability to detect different surface textures and complex patterns, and to detect object slip. The ability of current displays to meet human tactile thresholds for detecting contact, slip, pattern, vibration, and thermal sensations varies. The minimum bandwidth with which the human hand can perceive forces is 20-30 Hz and the majority of tactile displays meet this requirement. Pin-based displays are, in theory, capable of providing different patterns that can be sensed; since current devices have a pin tip of about 1 mm and pin spacing ranging from 1.5 to 3 mm, they meet or are close to the human thresholds for spatial resolution and the two-point limen. The maximum available pin array is, however, limited to 5 by 6 pins and this is insufficient for portraying any but the most simple patterns. In the case of vibration, displays seem to be evenly split between operating at low frequencies (<20 Hz) and mid-range frequencies (~200 Hz), this latter group being capable of presenting contact forces. While some experiments have shown the ability of at least one display to provide vibrations that support object manipulation, none approach the bandwidth recommended for supporting skillful manipulative tasks. In terms of frequency range, the largest range provided by any of the displays is 6 -100 Hz. Finally, the temperature displays are similar in providing a temperature resolution of 1°C which provides good support for the human JND. They vary in the range of temperatures that can be displayed with two out of the three displays capable of providing temperatures well in excess of what is likely to be needed, that is, temperatures beyond the human pain threshold.

Four tactile feedback products are commercially available. The *CyberTouch*, *Touchmaster*, and the *Tactools System* provide tactile displays that are mounted on the user's fingertips to provide feedback on object contact. The *Displaced Temperature Sensing System* generates thermal feedback, again via displays in contact with the user's fingertips.

Reflecting the immature status of this area, however, these products all are primarily display devices with a primitive software interface that requires the user to explicitly control the device. This first generation of products are best suited for use as research tools.

This is an relatively active area of research and much progress in addressing the issues outlined above should be made in the next few years. Even as research issues start to be resolved, practical problems in engineering and manufacturing small displays that can present tactile sensations to various hand and body areas may continue to limit practical use. With all these concerns in mind, it is unlikely that tactile feedback will come into widespread use in the next two to three years, though some initial practical use can be expected shortly thereafter. The switch to common use will be rapid, however, as soon as practical applications that demonstrate the value of tactile feedback appear.

As previously noted, the development of force feedback devices for use in VEs has greatly benefited from earlier work in providing force feedback for telerobotic applications. Accordingly, some parts of this interface technology are more mature than their tactile counterparts, although much progress is still needed.

The majority of current force feedback devices can be distinguished as exoskeleton devices that deliver forces to some subset of the shoulder, elbow, wrist, and finger joints; tool-based devices that deliver forces to the hand via a knob, joystick, or pen-like object held by the user; thimble-based devices that deliver forces to the user's fingertips; or robotic graphics systems that use real objects to provide forces to the hand. There are two exceptions to this categorization. Aura Systems, Inc. Interactor devices use low frequency sound vibrations to simulate force sensations that are presented to the user's torso, and the Rutgers Master delivers grasping forces to the hand via pneumatic micro-cylinders mounted on a glove. Exoskeleton devices have the advantage of allowing a user some freedom of movement in a VE, but are encumbering and their mechanical implementation may impose some restrictions on the joint movements. With the exception of the force feedback interfaces used in UNC's molecular docking and atomic force microscope, all the tool-based devices are desktop-based, thus constraining user movement. The desktop-based devices vary quite widely in the working space they support, ranging from only to few centimeters to a sphere of 40 cm diameter. The devices are primarily mechanical, driven by servo motor actuators. This technology present several problems, such as backdrivability and friction. The primary difficulty, however, is one of stability. The robotic graphics approach to providing force (and tactile) feedback is unencumbering and allows full user movement with a theoretically unlimited working space. Here the major issue is that of safety.

Among all these devices, eleven commercially available force feedback products have been identified. Since they are of very differing types and provide markedly different capabilities, each is suitable for different types of applications. Consequently, even ignoring performance characteristics, a prospective user is likely to have little choice among products. These systems are all expensive and most are developed to order, often with a significant delay before delivery. As yet, none of the available systems has seen significant

practical use. This situation is likely to change in the very near future because the newly released PHANTOM seems to be quickly becoming the system most commonly used by researchers.

In investigating how to evaluate the quality of force feedback systems, Rosenberg (1995) has proposed a set of minimum performance standards. While the necessary maximum force output and range of motion is application dependent, Rosenberg recommends a force output resolution of 12 bits, position resolution of 0.001 inch, and passive friction less than 1% of the maximum force output. Other requirements pertain to the system bandwidth (> 50 Hz), minimum sampling rate (2000 Hz), and latency (1 msec). Several systems meet some subset of these requirements, but currently only PHANTOM meets them all (with the possible exception of sampling rate, information on which was not available). Data on kinesthetic human capabilities collected through experiments provide other measures by which to assess force feedback systems. In this case, most current devices are capable of supporting the human JND for force sensing and the representation of a solid object to the fingers. Most are not, however, capable of providing the forces needed to represent a solid, immovable wall.

The hardware limitations of force feedback devices constrain the fidelity with which real world interactions can be simulated. In particular, the accuracy of sensors, latency of computer, performance of actuators, different location of sensors and actuators, and transparency of mechanical transmission all play an important role. Any force feedback device must allow for variability in hand size, otherwise the resulting scaling up or down of the force applied on the fingers will lead to imperfect perception of the interaction forces. Additionally, the characteristics of a user's interaction can change dynamically and radically, resulting in a non-linear system. Representation of rigid objects is a particular problem and most systems exhibit contact instability near a hard surface. Current approaches to this problem either add viscosity, which usually means that the user feels resistance even in free space, or reduce the stiffness of the simulated surface, leading to a spongy feeling. Since many aspects of this stability problem are insoluble, further understanding of how to employ multi-sensory input and how to exploit limitations in the human haptic system to alleviate the problem are needed. Safety is an example of another issue that needs further investigation. This is a concern that arises when there is a need to for the device to exert forces to oppose a user's volitional movements and safeguards are required to ensure that a computer or device malfunction does not result in user injury.

The software support required to implement force feedback interfaces is just starting to receive significant attention. Ad hoc force models to compute and generate forces have been developed by different researchers for use as research tools for quite some time. Now a few researchers are looking at looking at more general purpose model frameworks and developing techniques for more efficient haptic rendering. The work by Dr. Salisbury and his colleagues at MIT is notable in this area. Nonetheless, as force feedback devices continue to be developed, the lack of adequate software support remains a limiting factor in overall force feedback interface technology.

Another shortcoming lies in the understanding of human kinesthetics. While some investigation of human haptics with respect to the use of force feedback in VEs has been conducted, much more is needed. General issues in the areas of the biomechanical, sensorimotor, and cognitive abilities of the human kinesthetic system need to be investigated to provide better support for the hardware and software design of these interfaces.

As indicated above, current force feedback interface systems are severely limited in the types of force sensations they can deliver. Accordingly, within the next two to three years, the use of force feedback interface in practical applications is expected to be infrequent and mainly limited to a few application domains where the provision of force feedback is critical, such as surgery. Since many such applications will require special-purpose force feedback interface systems, new force feedback systems will continue to be developed. It is important to note, however, that although the current types of force feedback devices are likely to serve valuable roles in certain specialized applications, the approaches being taken are incapable on being scaled up to provide forces that more fully support the possible range of human interactions with a VE. Today, only the robotic graphics approach has the potential for such flexibility, and this is still the subject of feasibility studies.

7. FULL BODY MOTION INTERFACES

While full-body motion is commonly viewed as the most challenging VE interface technology to be developed, it is important to note that some types of full-body motion are feasible with current technology. Consider first those cases where a user is passively moved through a VE in a vehicle. Here, the usual practice is to build a "cabin" that represents the physical vehicle and its controls, mount this cabin on a motion platform, and generate virtual window displays and motion commands in response to the user's operation of the controls. These systems tend to be specialized to a particular application, for example, flight and tank simulators, and have been in use by the Department of Transportation, the Department of Defense, and the airline industry for many years. Indeed, cabin simulators represented the first practical VE applications.

Recent years have seen the exploitation of this technology by the entertainment industry for interactive VE adventure rides. Examples include IWERKS Entertainment's Loch Ness Expedition where, in a player-controlled submarine with periscope and robotic arms, six players try to save Nessie from bounty-hunters. Magic Edge, Inc. has developed a ride that sends twelve players, led by a squadron commander, on strike missions in X-21 hornets. In Galaxian-3, another Magic Edge adventure, players crew a star ship in a space battle. Greystone Technology, Inc. has developed the Mercury VR Platform, a futuristic flying motorcycle used by players to participate in the MagBall team game, using simulated magnetic fields generated by their craft to manipulate a ball and score goals. Other Mercury rides include Canyon Runner, a game where players participate in a futuristic Gauntlet League race using guns to eliminate rival competitors and simulated kinetic fields to deflect enemy shots and beams from canyon-mounted pulse cannons. Chameleon Technologies, Inc. use a centrifuge-based system with cabins, suspended from up to ten arms, capable of full 360° movement. Three Chameleon games are currently available, a futuristic space game called Labyrinth Rangers, a drive-and-shoot race car game called LazerDrive, and the MERCS supersonic aircraft mercenary game; players continually interact with each game, for example, executing aircraft barrel rolls and dives in accordance with the game objectives.

For many kinds of VE applications, however, more active self-motion is required. With the limiting constraint of a stationary surface under the user that naturally provides all necessary kinesthetic cues, simple in-place user movements in a VE only require the generation of appropriate visual displays. If the surface is uniform but moving, a motion platform can be used to provide the necessary motion cues. Even locomotion through a *small*

(typically around 10 x 10 feet) virtual space poses no significant problems, as long as there is a surface that can provide the necessary kinesthetic cues. The major challenges for full-body motion in a VE arise whenever any of the following are required: locomotion through a large virtual space, locomotion over varying surface characteristics, and motion in a direction other than horizontal.

This section starts with a brief look at the relevant human sensory capabilities. The following two sections deal with interfaces that, respectively, support active and passive motion through a VE. The final part of this section presents expectations for the development of full-body motion technology in the next five years.

7.1 The Human Motion Sense

Many systems play a role in a human's capability to sense motion and control posture (orientation and balance), the two primary systems being the visual and vestibular systems. Some details about the visual system have already been presented in Section 2.1. In the context of motion, however, it is important to note that the visual system is both a sensory and motor system. In the former case, it signals the position and movement of the head with respect to surrounding objects, and provides information about the direction of the vertical. As a motor system, the visual receptors that sense slipping of the retinal image supplement compensatory eye movements through a tracking mechanism called the optokinetic reflex.

The vestibular system also is both a sensor system and a motor system. In its role as a sensory system, the vestibular system provides information about movement of the head and the position of the head with respect to gravity and any other acting inertial forces. It uses two types of sensory organs. The first of these are the semi-circular canals in the inner ear that provide information about the angular velocity of head movements. These canals are fluid-filled and the inertia of this fluid causes head rotations to increase, or decrease, activity of specialized hair cells that fire neural signals to excite the vestibular nerve. The neural firing in the vestibular nerve is proportional to head velocity over the range of frequencies in which the head commonly moves, that is, 0.5 to 7 Hz. However, the semi-circular canals provide the best response in the first second or so, and output decays exponentially with a time constant of about 7 sec. The set of three canals on each side of the head work in a complimentary push-pull relationship, with the canals in each set being aligned perpendicularly to each other. This alignment allows the two vertical canals to signal forward and backward head rotations, while the horizontal canal signals rotations about the vertical axis. The second type of vestibular sensory organ is the otolith organ. There are two otolith organs associated with each set of semi-circular canals and they provide information about linear acceleration and head tilt with respect to the gravitational axis. The saccular otolith provides information about vertical linear acceleration of the head, and the utricular otolith responds to horizontal accelerations. There can be ambiguity in, for example, determining whether an anterior head rotation signalled by the semi-circular canals was the result of head flexing on the neck or body flexing at the waist. Signals from the visual

and other systems are used to resolve these ambiguities when they occur. In general, the semi-circular canals respond best to rapid head movements, while the otoliths are most sensitive to slow movements.

As a motor system, the vestibular system plays an important role in posture control, that is, orienting to the vertical, controlling center of mass, and stabilizing the head. To this end, output from the vestibular system goes to the spinal cord to serve the vestibulo-spinal reflex. This reflex generates compensatory body movements to maintain head and postural stability. Output from the vestibular system also goes to the ocular muscles serving, in this case, the vestibular-ocular reflex that generates eye movements that enable clear vision while the head is in motion.

Benson (1990) has summarized the findings of several researchers on the functional thresholds of the vestibular system. He reports that, using a seat free to move in the x or y body axis, the threshold for detection of tilt from the vertical is on the order of 2° . The perception of angular motion varies with frequency, falling at around 0.2 log unit/decade between 0.1 and 1.0 Hz, and falling at -1 log unit/decade below 0.1 Hz. For stimuli shorter than 15 seconds, this perception of angular motion is related to the time, t , taken to detect angular acceleration, α ; the product αt has a mean constant value of $3.7^\circ/\text{sec}$. For sustained rotational stimulation with prolonged acceleration (such as can occur in an aircraft), the sensory threshold for angular rotation is determined by the magnitude of angular acceleration rather than velocity change and the mean threshold for angular accelerations of the head about the z axis has been demonstrated as $0.32^\circ/\text{sec}$ with a range of 0.05 to $2.2^\circ/\text{sec}$. With respect to the perception of linear acceleration, for a linear oscillation at approximately 0.3 Hz in the horizontal plane, the mean threshold was around 0.03 m/sec^2 for oscillations in the x, y axes and around 0.06 m/sec^2 for oscillations in the z body axis. The common peak angular velocity for passive nodding of the head, such as occurs during walking or running, is $\pm 10^\circ/\text{sec}$. Volitional head movements usually exhibit a peak angular velocity of at least $100^\circ/\text{sec}$ but may be as high as $500^\circ/\text{sec}$. Peters (1969) summarizes various experimental findings on the threshold for detection of motion about the vertical axis, reporting that the threshold ranged from 0.2 to $2^\circ/\text{sec}^2$. The threshold for linear acceleration has been found to range from 0.002 to 0.027 g.

There are circumstances in which other sensory systems impact the sensory thresholds of the vestibular system. For example, Huang and Young (1981) found that while the level of illumination produces no significant differences in the threshold for perception of angular velocity, the absence of illumination significantly lowers the threshold and reduces latency time.

Benson describes several functional limitations suffered by the vestibular system. Transient movements lasting less than 10 sec with a change in angular velocity below roughly $2^\circ/\text{sec}$, or peak acceleration below roughly 0.05 m/sec^2 , may be undetected. Prolonged rotation of the head (over about 15 sec) with cross-coupled stimulation of the semi-circular canals can cause misperceptions. Misperceptions of altitude can occur in the pres-

ence of prolonged (40 to 60 sec) linear acceleration, or deceleration, when the resultant effect of the imposed acceleration and head orientation is unaligned with the gravitational vertical. Head movements during linear accelerations over 10 m/sec^2 (1 g) also cause misperceptions of the direction of the movements, and when the acceleration increases to more than 50 m/sec^2 , head movement can cause the perception of tumbling.

The different forms of illusionary passive self-motion have been studied for many years. Such perceptions can be generated by vestibular stimulation, for example, by sinusoidal stimulation of the horizontal semi-circular canals, stimulation of the cervical neck receptors, or visual stimulation. In general, visual and cervical stimulation dominate vestibular stimulation. Since only visual stimulation is likely to be used in VEs, the rest of this discussion is so restricted.

Linearvection, the illusion of linear motion in a stationary individual, is known to be generated by moving images in the visual field. In a series of experiments, Berthoz, Pavard, and Young (1975) measured image velocity and luminance thresholds for the appearance of linearvection. The thresholds of differential luminance level decreased with increases in image velocity, reaching a minimum level between 0.001 and 0.0001 cd/m^2 . The thresholds of image velocity differed depending on whether the moving image was presented only to the periphery of the visual field or the entire visual field. When inducing the sensation of forward self-motion, linearvection appeared at an image velocity of approximately 0.03 m/sec in the first case and approximately 0.01 sec in the second. These figures were substantially less when inducing backwards linearvection. The velocity of the perceived linearvection increased with the velocity of the image display, reaching a saturation point when the image moved at a rate of about 1 m/sec . The latency of onset for linearvection was around 1 sec . The researchers also investigated the effects of prolonged exposure to linearvection. Here they found that the time constant of adaptation to linearvection ranged between 30 to 50 sec , after which time subjects were prone to underestimate the velocity of induced motion. Finally, in the presence of conflicting visual and vestibular cues, Berthoz, Pavard, and Young found a dominance of visual cues. More specifically, it seems that the vestibular cues dominate the short-term subjective determination of acceleration, whereas the visual cues dominate in the long-term sensation of velocity.

Visual cues can also generate the illusion of circular motion, called circularvection. Huang and Young (1981) found that the perception of self-motion is significantly more sensitive when viewing an isolated visual target that is rotating with the subject, than in the absence of a visual target. Duration of apparent motion is usually longest and perception thresholds lowest under conditions of dim illumination and plain background. In reviewing the findings of other researchers on the thresholds for perception of angular velocity, Huang and Young report no consensus on the velocity threshold for perception of vertical angular acceleration in an unilluminated environment, with experimental findings ranging from 2.0 to $16.4^\circ/\text{sec}$. Significant interaction for the duration of induced self-motion between the simulated speed of the observer and the visual angle of the display has been observed by other researchers. The duration of reported self-motion was smallest for the largest visual

angle examined (21.1°) and the high speed condition. Highest depth ratings occurred in conditions in which the longest duration of self-motion was reported, possibly indicating that induced self-motion in the central visual field is dependent on relative depth information within the display. The latency of onset of circularvection is cited as ranging from 1 to 14 sec (Brandt, 1973).

Neck receptors are capable of inducing strong circularvection. (These receptors are usually stimulated by seating a subject in a chair inside a rotating drum, with the subject's head fixed in a clamp attached to the ceiling of the drum.) Both visually and cervically induced illusions of head rotation overrule the vestibular sensation of head movement when estimating head position (Bles and de Jong, 1982). Relative to the vestibular induced sensation, not only the visual but also the cervically induced sensation of head motion is strong. For stimulation of the vestibular, visual, or cervical systems separately, the size of actual head movement is generally underestimated. Pure visual stimulation can fail to induce circularvection, although combining vestibular plus visual stimulation has indicated that vision, and not the vestibular system, determines circularvection.

7.2 Self-Motion Interfaces

Self-motion interfaces are defined as those cases where the user moves himself through a VE, as opposed to being passively moved in some type of vehicle. Currently the illusion of self-motion through a VE is supported by generating visual displays that represent some concept of "flying" when the user points a finger or some type of wand in the direction he wishes to travel. Undoubtedly, there are many types of application for which such interaction is ideal, but flying through an environment may well give a different perspective and less detailed knowledge of the environment than that which can be acquired by walking through it. In particular, locomotion is needed to acquire accurate information about surface characteristics such as resilience, slope, and texture. It is also essential for time-related information, in those cases where visibility is limited in some fashion, or when a user is required to exert the types of energy he would in performing actions in the physical world. These differences will be critical when VEs are used in applications such as special operation forces mission rehearsals.

7.2.1 Commercial Products

The interface mechanisms commonly used to control VE flying are the whole hand input and pointing devices discussed in Section 5.1 and Section 5.2; these are not mentioned further here. Likewise, exercise systems that link conventional exercise equipment with the 2-D presentation of scenery on a CRT are not considered representative of VE systems and also excluded. Instead, this section discusses a range of novel types of interface devices (such as gyroscopes, hang gliders, and interactive motion platforms) that only recently have become commercially available and that provide kinesthetic motion feedback.

Human gyroscopes allow a user to rotate his body axis freely in 6 DOFs and are available from Aerotrim USA, Inc., Orbotron, Inc., RPI Entertainment, and StrayLight Corporation. With appropriate position tracking of body movements, these devices allow other interface displays to be coordinated with the user's movement. Virtual Images' CyberPak allows the user to turn freely in any horizontal direction while he controls his rate of forward or backward motion through a hand-held controller. A hang glider is available from Dreamality Technologies, Inc. and Trailcraft Manufacturing Ltd. This device allows a representation of gliding through a VE, with the direction of motion controlled by the user turning his body and pushing on some type of bar. Information on a second hang glider, marketed by CyberEvent Group, Inc., was not available. A motion platform manufacturer, Denne Development Ltd. (DDL) has developed products that offer new ways of full-body interaction. By making the motion platform itself responsive to changes in the user's center of gravity, and providing information on these changes to the VE system, these interactive platforms also allow coordination with other interface displays. The final product discussed is a very different type of interactive motion base system developed by RPI Entertainment. The characteristics of these various products are summarized in Table 18.

Table 18. Characteristics of Commercially Available Self-Motion Products

Product	Vendor	Device Type	Range of Motion	Additional Provided Interface Equipment	Price
Aerotrim	Aerotrim USA, Inc.	Gyroscope	360° pitch, roll, yaw	None	\$7,995
CyberPak	Virtual Images, Inc.	Revolving backpack	360° horizontal	HMD, head tracker, 3-D localized sound display	\$31,500
CyberTron	StrayLight Corporation	Gyroscope	360° pitch, roll, yaw	HMD, head tracker, sound display	\$54,000
DreamGlider	Dreamality Technologies, Inc. and Trailcraft Manufacturing Ltd.	Hang glider	~60° horizontal	HMD, sound display	\$26,700
Orbotron	Orbotron, Inc.	Gyroscope	360° pitch, roll, yaw	None	\$10,000
X-otron VR	Orbotron, Inc.	Gyroscope	360° pitch, roll, yaw	HMD, head or ring tracker, 3-D localized sound display, hand controllers	\$27,000 ^a
Supertron	Orbotron, Inc.	Gyroscope (2-man)	360° pitch, roll, yaw	None	\$14,500
PemRAM 3 Axis Motion Base	Denne Developments, Ltd.	Interactive motion platform	42° pitch, 40° roll, 360mm heave	None	\$27,000
PemRAM 6 Axis Motion Base	Denne Developments, Ltd.	Interactive motion platform	40° pitch, 40° roll, 50° yaw, 0.4m heave, 0.4m sway, 0.45m surge	None	\$50,000
SimuPod	RPI Entertainment	Rotational	360° pitch, roll, yaw	Visual display (various), 3-D localized sound display, hand controllers, "rumble and thump" generators	From \$30,000
SimuSled	RPI Entertainment	Forward tilting motion platform	100-360° pitch, 10-45° roll, (optional 360° yaw, shift motion axis)	Visual display (various), voice recognition, 3-D localized sound display, hand controllers, "rumble and thump" generators	\$14,000

a. The motorized, interactive version starts at \$35,000.

7.2.1.1 Aerotrim

Aerotrim USA, Inc. developed their patented gyroscope motion platform, Aerotrim, for use as an exercise machine, although it has since been used for a variety of purposes ranging from aiding pilots to overcome airsickness and disorientation to therapeutic treatment for neurological disorders. It has also been used by Sportsland America and others as a means of user navigation in VEs.

The Aerotrim is a free-standing gyroscope that allows the user to control whole body orientation in any direction. A photograph and some specification details are given in Figure 103. The Aerotrim base price is \$7,995.




Figure 103. Aerotrim

7.2.1.2 CyberPak

CyberPak from Virtual Images, Inc. is another interface system designed to allow a user to navigate through a VE. The user stands on a stationary platform, positioned against a NASA-like space pack, and holds a hand grip and control buttons attached to the pack via arm structures. He is free to turn in any horizontal direction and controls forward and backward motion through the VE by pushing or pulling on the hand grip. The system is controlled by a Pentium PC with a Division image generator. A choice of HMDs is available, supported by head tracking. CyberPak also includes 3-D binaural sound spacing for its audio interface. Figure 104 provides some specification details.

Specification	
Floor Space	6 x 9 ft
Device Height	7 ft
User Range of Motion	360° horiz
Device Weight	~400 lb
Max. Payload	> 1,000 lb
User Size	19 in width

Figure 104. CyberPak

The system primarily is intended for use in entertainment applications and comes with a TV monitor for external audience viewing and one game (additional games will shortly be added). Virtual Images, Inc. also see CyberPak being used in diverse applications ranging from training to therapy for fear of heights. Currently only supporting individual users, a networking

capability is expected to become available in Spring '96. The price for CyberPak starts at \$31,500.

7.2.1.3 CyberTron

The CyberTron system, developed by StrayLight Corporation, provides a VE interface that combines visual (using a Liquid Image HMD) and auditory feedback with user motion. Motion is supported by a gyroscope motion platform, chiefly used so that the user can simulate flying through a VE. A Polhemus Isotrak II is used for head tracking so that user movements can be monitored and used in the generation of visual displays. A photograph of the CyberTron and further details are given in Figure 105.





Figure 105. CyberTron

StrayLight Corporation has developed three standard games for use with CyberTron, including one game that allows competitive play between users in up to four networked CyberTron systems. CyberTron, including HMD, tracker, and one game is priced at \$54,000.

7.2.1.4 DreamGlider

The DreamGlider system from Dreamality Technologies, Inc. and Trailcraft Manufacturing Ltd. provides the experience of hang gliding through a VE. It comprises a trapeze support system that provides the sensation of vertical motion to the user, who rests in a supporting sling and controls the direction of the glider by forces he exerts on a control bar. The system runs on a Pentium-based PC, augmented with a Synthetic Images Reality Blazer image generation board and an Advanced Gravis Ultrasound card. Currently, it uses a modified Virtual I/O HMD, although Dreamality Technologies in collaboration with Forte Technologies is developing their own HMD that is expected to be introduced later in '96. A photograph and some specification details for the system are given in Figure 106. The price for a single DreamGlider system, including one hang gliding game, is \$26,000.

Currently, three special effects are being developed to augment the VE experience. The first of these is a smell capability, expected to allow the release of up to four different scents as



Specification	
Floor Space	7 x 8ft
Device Height	7 ft
User Range of Motion	Side-to-side, forward-back, and sway movements typical in hand gliding
Device Weight	500 lb
Max. Payload	325 lb
User Size	N/A

Photo courtesy of Dreamality Technologies, Inc. and Trailcraft Manufacturing, Ltd.

Figure 106. DreamGlider

a user flies through certain areas in the VE. The second effect will be the addition of a hydraulic servo motor capable of raising and lowering the DreamGlider device to provide the sensation of lift. Finally, fans will be included to generate the sensation of a breeze.

Originally designed as an entertainment system, the DreamGlider system is being upgraded for use as a hang gliding trainer. The current system starts the hang gliding experience with the user stepping off a cliff. For training applications, the user needs to be able to run down a hill and then take off and the companies are investigating how a treadmill can be integrated with the DreamGlider device to introduce this capability.

7.2.1.5 Orbotron, X-otron VR, and Supertron

Orbotron, Inc. market a series of human gyroscope motion devices. The initial product, also called Orbotron, is not motorized and allows free motion in pitch, roll, and yaw dimensions for a single user. Initially designed as an workout machine, the Orbitron was quickly used for entertainment applications. With the VE option, called the X-otron VR, it includes an Optics 1 HMD capable of either 2-D or 3-D optics, with either head tracking, or user tracking based on the position of the gyroscope rings. The VE system also includes localized 3-D sound generation and supports dual joysticks with a total of eighteen programmable buttons for user input. Further details for the X-otron VR are given in Figure 107. The price of the Orbitron is \$10,000, and the X-otron costs \$27,000.

Specification	
Floor Space	10 x 10 ft
Device Height	9 ft 8 in
User Range of Motion	360° pitch, roll, yaw
Device Weight	1,670 lb
Max. Payload	300 lb
User Size	4 to 6.16 ft or 4 to 6.33 ft

Figure 107. X-otron VR

Specification	
Floor Space	114 sq ft
Device Height	11 ft
User Range of Motion	360° pitch, roll, yaw
Device Weight	1,950 lb
Max. Payload	300 lb
User Size	4 to 6.16 ft or 4 to 6.33 ft

Figure 108. Supertron

The more recently introduced Supertron differs from the Orbitron in allowing free motion for two users simultaneously. This device was designed for either entertainment or research applications. Currently available as a stand-alone device, priced at \$14,500, a Supertron VE option is expected to become available by the end of 1996. Further details for the Supertron are given in Figure 108.

A motorized single-man gyroscope that can be controlled by a computer, or an analog input from some other source, is also available. Designed as an entertainment device, this product supports VE applications with the same HMD, sound display, and joysticks as the Orbitron, and both head tracking and ring tracking are supported. In appearance similar to the Orbitron, the motorized single-man gyroscope differs in requiring 12 x 12 feet area of floor space. Its price starts at \$35,000. A two-man version of the motorized gyroscope is expected to become available in Fall '96.


Wheelchair versions of both the Orbitron and the single-man motorized gyroscope are available. These are suitable for use by paraplegics or quadriplegics.

7.2.1.6 PemRAM Motion Bases

Denne Developments, Ltd. (DDL) market new motion platforms based on their patented Precision electromagnetic RAM (PemRAM) actuators. In these electromagnetic rams the space under the pistons is filled with air at a pressure sufficient to support the deadweight of the payload. The pressurized part of the ram is connected to a small reservoir and isolated from the main air supply so that it forms a long-stroke gas spring. This counterbalance system enables the motion base to stay where it is, effectively in neutral equilibrium, and the dynamic motion is provided by impulsive forces from the electromagnetic actuators. The force actuator generates forces that are felt as acceleration cues by the user. By effectively eliminating gravity from the equations of the motion, the actuators have a low power requirements, typically one tenth of that required for hydraulic systems. The high bandwidth of the system allows quick and precise control of the forces generated, including vibrations exceeding 30 Hz.

A feature that makes the new PemRAM motion bases potentially very useful for VEs is their ability to automatically react to user movements. Since the currents flowing in the actuators are continuously monitored, it is possible to identify changes in the center of gravity of the platform, that is, identify the movements of the user and allow this to be taken into consideration in interactively controlling the motion base and any other VE interface displays. When the user leans in any direction, the actuators automatically adjust to compensate for the shift in the center of mass, holding the position of the motion base constant.

Two PemRAM motion bases are currently available, a three axis and a six axis system. Via an electronic control unit, these motion bases provide a serial computer interface. Further details are given in Figure 109 and Figure 110. The prices are approximately \$27,000 and \$50,000 for the three axis and six axis motion bases, respectively.



Specification			
Base	1.5 m equilateral triangle		
Weight	240 Kg		
Payload	300 Kg nominal, 500 Kg maximum		
Loading Height	590 mm		
Capsule Interface	Mounting triangle on 620 mm radius pitch circle		
Performance:	Displ.	Velocity	Accn.
Heave	± 180 mm	0.5 m/sec	0.5 g
Pitch	$-24^{\circ} + 18^{\circ}$	35°/sec	3330°/sec ²
Roll	$\pm 20^{\circ}$	40°/sec	330°/sec ²

Photo courtesy of Denne Developments, Ltd.

Figure 109. PemRAM 3 Axis Motion Base

Specification			
Base	1.5 m circle		
Weight	380 Kg		
Payload	600 Kg nominal, 1000 Kg maximum		
Loading Height	0.6 m		
Capsule Interface	Mounting triangle on 0.7 m radius pitch circle		
Performance:	Displ.	Velocity	Accn.
Heave	± 0.2 m	0.5 m/sec	0.5 g
Pitch	$\pm 20^{\circ}$	35°/sec	3330°/sec ²
Roll	$\pm 20^{\circ}$	40°/sec	330°/sec ²
Yaw	$\pm 25^{\circ}$	40°/sec	330°/sec ²
Surge	+0.25m, - 0.2m	0.5 m/sec	0.5 g
Sway	± 0.2 m	0.5 m/sec	0.5 g

Figure 110. PemRAM 6 Axis Motion Base




Photo courtesy of Denne Developments, Ltd.

Figure 111. PemRAM Surfing Demonstration System

The interactive capability of the PemRAM motion bases has been demonstrated in a number of entertainment applications, including one for surfing (see Figure 111) and another for flying through an aerial obstacle course. These applications have used a flat screen rather than an HMD because of the disparity between the high update rate for PemRAM motion cues (ideally 100 Hz) and that typical for commercially-available HMDs. Also, the demonstrators have found that users able to maintain a connection with the real world through their peripheral vision fail to experience the disorientation that has occasionally been reported when peripheral vision is not provided. Additionally, for safety reasons, these demonstrations have surrounded the user with a hand rail mounted on a platform.

7.2.1.7 SimuPod

RPI Entertainment's SimuPod is a rotational motion-based product with a V-brace configuration linear actuator. It allows a full 360° movement for pitch, roll, and yaw, supporting any rotational body effect. Motion can result from the movements of the user's body, or be generated by a PC-based motion control system. Rumble and thump effects are produced using bass speakers and vibration transducers attached to the user's backpack. A variety of RPI developed visual displays are available for use with the SimuPod, including HMDs, projection, pull-up head-coupled, and lean-in head-coupled displays, and 3-D localized sound is available. For user input, hand controllers can be attached to the device frame. The SimuPod is available in wireless mode, where no cables are attached to the user. A photograph of the SimuPod and further details are given in Figure 112.




Figure 112. SimuPod

The SimuPod is primarily intended for the entertainment industry, and over 44 licensable rides and games are available. However, the product is available to developers of other types of VE applications. It is built and customized to order allowing, for example, restrictions on user size and weight to be adjusted to meet client needs. The price for SimuPod ranges from \$30,000 to \$350,000.

7.2.1.8 SimuSled

Another RPI Entertainment product is the SimuSled. The user starts by standing up and then falls forward or backward, allowing simulation of flying, skydiving, gliding, and sledding, or any similar activities. The basic product includes a motion platform capable of 100° pitch and 10° roll motions, though these can be increased to 360° pitch and 45° roll. An optional 360° yaw and a fourth motion axis, described as shift, can be added. The SimuSled includes a voice input channel to accept user spoken commands. As with other RPI Entertainment products, a range of visual displays, 3-D localized sound, and special effects are available and the device is network capable.

As before, this product is primarily intended for the entertainment industry, and many games are available, but it can also be purchased for use in other types of applications. It is built and customized to order, with an approximate price of \$14,000 for a single sled, depending on configuration.

7.2.2 Current Research and Development

Until recently, research into self-motion interfaces was exclusively the province of universities. Over the past few years, the DoD has started to support research in this area, leading to the involvement of some private companies. While the DoD is focused on providing support for dismounted infantry actions in the military's Distributed Interactive Simulation (DIS) network, it is likely that the bulk of the research products will be applicable to the more general, non-military VE self-motion applications. Unlike the university work, however, much of the DoD-sponsored work is short-term; for example, the Small Business Innovation Research (SBIR) efforts discussed below are all required to be completed within a six month interval. Whether the DoD will continue to support research in this area is unknown.

7.2.2.1 Computer Graphics Systems Development Corporation

Researchers at CGSD are investigating the development of a Locomotion Simulator for Three-Dimensional Virtual Space. The goal is to produce a 3-D generalization of a treadmill, called an OmniTrek, that will enable walking and running in virtual space without actually traveling more than four feet from a nominal position. This device is intended to support turning to any direction and simulation of climbing or descending stairs, or travelling across variable terrain. It is designed to be used safely without a tether or harness.

The OmniTrek is a computer-controlled robotic apparatus about ten feet in diameter. The top surface of the device is co-planar with a raised floor. The raised floor is about four feet above the primary floor upon which the device sits. The walking surface has no holes or gaps, and is flat except when climbing or descending. Simulation of some soil types, such as sandy soil, is possible. The device is designed to output the position and velocity of the user, and in most cases the position of the user's feet and head. The most difficult aspects of the design relate to the safety of the device, to ensure that the user will not easily lose balance in ordinary operation and that, in any case, a fall would not result in a serious injury. The general approach is to use three levels of safety: software that controls the device to prevent injury, hardware detectors that stop motion if a foot is in a potentially dangerous place, and design details that will push the user out of the way if all else fails. The virtual reality imaging system used with the OmniTrek must have low latency and generally would include graphic imagery of the body of the user. The preferred computer interface will be dedicated Ethernet.

The OmniTrek is the subject of a SBIR feasibility and design study sponsored by the US Army Simulation Training and Instrumentation Command (STRICOM). The design and the study are currently nearing completion, and demonstrations of certain of the control aspects (control laws, tracking problems, and motor sizing) have been made in the context of a one-

directional device. If the feasibility and design study is successful, the researchers hope to develop a working prototype system that would include imagery for the virtual environment.

7.2.2.2 Cybernet Systems Corporation

In another effort sponsored by STRICOM through a SBIR contract, Cybernet Systems Corporation is developing a prototype full body kinesthetic display to simulate locomotion for dismounted troops under virtual training and exercise scenarios. The design is referred to as the foot haptic approach because it begins by providing a full six axis motion platform for each foot of the soldier. To this base design, kneeling boards are added to support rolling, kneeling, and prone postures. A vertical feature presentation mechanism allows pushing realistic walls, windows, doors, and high vertical obstacle features.

Commercially, the system will be applicable to a variety of motion-based training and play scenarios. It can be the basis of systems for indoor track and field training, and training in eye-hand-body coordination sports such as tennis, baseball, and golf. The system can be incorporated into training VE systems for civilian safety and police personnel. Finally, it can be used in VE-based rehabilitation systems.

7.2.2.3 Institute for Simulation and Training

Dr. Jim Parsons at the Visual Systems Laboratory of the Institute for Simulation and Training (affiliated with the University for Central Florida) is developing a treadmill-based locomotion system. The treadmill being used is of a type used by cardiologists in assessing patients' medical conditions. It has been modified in several ways. The motor was removed to increase the safety of the device for users wearing HMDs and, to increase ease of user movement, the plywood foundation under the belt was replaced with a specially made material-handling conveyor. User motion is detected via proximity switches mounted on each edge of the treadmill and the user signals turns to the left or right using buttons mounted on the treadmill handles. A photograph of the device is shown in Figure 113.



Figure 113. IST Treadmill Locomotion Device

The signals from the treadmill to the Silicon Graphics machine that generates the VE visual scenes are transmitted via a modified Microsoft Mouse: the signals from the directional buttons on the treadmill are wired into the mouse buttons, and the signals from the proximity switches drive relays that simulate the roller ball of the mouse. The serial mouse input is then translated to provide the proper granularity for the VE, allowing the rate of user motion to adjust the visual display.

This treadmill locomotion device is being used by the Army Research Institute in a series of experiments that are investigating the effectiveness of the VE training for outdoor nav-

igation skills. Current work with the interface system is focusing on providing the ability to monitor finer granularity of user motion.

7.2.2.4 Sarcos Research Corporation

Sarcos performed some of the earliest work in linking the individual combatant to VEs. Initial work in the Individual Portal into Virtual Reality (IPORT) series was conducted under contract for Army Research Laboratory (ARL) to build a UNIPORT motion platform. This device (shown in Figure 114) allows the user to direct his movement through a VE by turning the seat to change his direction and adjust speed of motion by the rate of pedaling; the visual scenes generated are tied to these motions. An important feature of the UNIPORT is that it allows the physical exertion imposed by moving through the VE to be adjusted based on such characteristics as the type of motion being simulated (walking, crawling, running), the terrain surface, and the load the soldier is expected to be carrying. The UNIPORT provided the first linkage of physical exertion into the VE and the first integration of dismounted infantry in the Distributed Interactive Simulation network. When linked into military simulations conducted on the DIS, the UNIPORT's capability to require appropriate action-related physical exertion provides for soldier decision-making under conditions of risk.



Photo Courtesy of Sarcos Research Corporation

Figure 114. UNIPORT



Photo Courtesy of Sarcos Research Corporation

Figure 115. TREADPORT

Building on this initial work and an analysis of infantry movements and combat actions, SARCOS completed a functional requirements description for linking Dismounted Infantry to VEs, and the Systems Architecture is in final draft. This work was sponsored by ARPA in support of the Army's Dismounted Infantry Battlespace BattleLab (DIBBL).

SARCOS' second development in the IPORT series, the TREADPORT, allows the soldier to move naturally within the VE. A tether and control system keep the soldier essentially

centered on the flat treadmill-like surface of the TREADPORT and the VE is adjusted responsively to his movements. It has a natural user interface allowing walking/running/crawling and a full range of postures—kneeling, sitting, prone, etc. A photograph of the TREADPORT is shown in Figure 115. This work was sponsored by ARPA and STRICOM in support of the DIB-BL.

Nearing completion is the Individual Soldier Mobility Simulator (ISMS) being developed for the ARL. This motion platform (shown in Figure 116) was developed as a high fidelity terrain interactive system including feedback for mud, rocks, stairways, and navigational obstacles. The system will support testing of new soldier equipment in a laboratory environment representative of field conditions. It features scenario-based energy expenditure and high fidelity foot-motion tracking.

All SARCOS motion platforms will work with Polhemus or the SARCOS SENSUIT for motion capture. The SENSUIT allows direct, interactive, real-time control of an icon in a virtual world. It is a clear and responsive icon-wise to the hand and arm signals required for dismounted infantry combatant operations. Compared to the Polhemus, SARCOS states that the SENSUIT is free-ranging, features greater accuracy and speed in measuring joint angles directly, provides better resistance to sensor drift, has the capability to measure many more degrees of freedom, does not require inverse kinematic computations, and is insensitive to metals or other interference in its environment.

7.2.2.5 Systran Corporation

In the final SBIR I effort being sponsored by STRICOM, Systran Corporation also is completing a feasibility study for a locomotion simulator to support training of dismounted infantry in the DIS environment. This work has centered on the development of a design for a simulator called LocoSim.

In the simplest terms, LocoSIM employs DDL's PemRAM actuators (see Section 7.2.1.6) to move boot plates to various positions within an operating envelope defined by the granularity of the micro-terrain being traversed. The user still moves as usual in walking, running, or crawling (based on foot movements only) over different terrain, or ascending or descending stairs. His leg and foot actions initiate actuator displacements that, via the connection of linkages and size of displacements, result in the displacement and angular orientation of each boot plate. A pivot plate attached to the forward end of the boot plate allows the user to change the direction of his motion and provides a limited rolling capability while in a prone position. Each boot plate is individually controlled, with the front and back being driven separately to allow matching the characteristics of human gait. An additional actuator is used



Photo Courtesy of Sarcos Research Corporation

Figure 116. Individual Soldier Mobility Simulator

to provide a side-step capability. A leg/foot sensor suite (including force, displacement, and rotary sensors) is mounted on an exoskeleton, itself mounted on the boot plate, to provide input on user movements. The LocoSim itself is intended to be placed below floor level, underneath a gap in the floor, such that the boot plates are positioned horizontally with the floor. A sketch of the device is shown in Figure 117, while Figure 118 shows a more detailed view of the boot plate and sensor mechanism. As designed, each boot plate exhibits a DOF for each of the longitudinal, height, pitch, roll, and transverse axes. Freedom for the yaw axis is simulated via force-feedback from the boot plate and leg/foot sensor system. (DOFs in the roll and transverse directions are not expected to be supported in the initial LocoSim implementation).




Figure 117. LocoSim




Figure 118. LocoSim Boot Attachments

The overall LocoSim design includes the requirements for the LocoSim control system, intended to interface with the DIVE VE supported by the Army Research Laboratory, and a high level functional software design.

7.2.2.6 University College London, UK

Researchers at the University College London and the London Parallel Applications Center propose a paradigm of body-centered interaction as an alternative to traditional mouse and menu interaction schemes for user interaction with VEs. In accordance with this paradigm, Dr. Mel Slater's group has developed the Virtual Treadmill system. With this system, a participant can walk normally within the range of the tracking device. For moving over larger distances, the participant employs gestures similar to walking, in effect, walking in place. These gestures are identified by pattern analysis of head tracking data using neural nets. In addition to walking, the Virtual Treadmill allows a participant to walk up steps by detecting an object collision of the participant with a virtual bottom step, and then subsequent walking moves the participant up the steps. Walking down steps is initiated by identifying when the participant steps over the top step. Climbing and descending virtual ladders is supported in a similar manner, with the participant's hand elevated over his head to signal going up the ladder, and hand positioned lower than his head to signal descending the ladder.

This body-centered paradigm is assumed to help participants identify with the representation of their body in the VE, that is, their virtual body, and so increase their sense of presence. In an experiment to test this hypothesis, the researchers compared two different navigation schemes with respect to their effect on the reported sense of presence (Slater, 1994). Sixteen subjects were selected, half of whom used a 3-D mouse for navigation, and the other half used the Virtual Treadmill. The experimental task was to pick up an object, take it into a room, and set it on a particular chair. The chair was positioned so that a subject had to cross a 20-foot deep chasm to reach it, by either following a wide ledge around the room or going directly across the chasm. In all cases, the participant saw a virtual body as self representation, or at least those parts of the body that would normally be visible given the current direction of gaze. The main finding of the experiment, as reported by Slater, was that while a higher association with the virtual body gave a higher presence score for the Virtual Treadmill participants, there was no such correlation for the 3-D mouse users.

Current research by the group is largely concerned with multi-participant VEs where the participants may be at different physical, remote locations and the shared virtual world distributed over a wide-area network. Such a system is already linking the Universities of London, Nottingham, and Lancaster. In this case, not only is there a relationship of individual participants to their own bodies, but recognition of and interaction with the other participants. At the present time, this system uses body-centered interactions, based on head and hand tracking data, to allow gestural movements to be translated into 3-D geometrical, virtual structures. The next phase of development will include the integration of the Virtual Treadmill into the system.

Future plans with respect to the Virtual Treadmill include modifying the system to recognize different walking speeds.

7.2.2.7 University of Tsukuba, Japan

For several years, Dr. Hiroo Iwata at the University of Tsukuba, Institute of Engineering Mechanics, has led researchers in the development of a Haptic Walkthrough Simulator. As the name suggests, the device is primarily intended to support walkthroughs of building and urban space designs.

The user of this device wears a pair of modified roller skates that are equipped with four castors to permit him to move in any direction. Each skate has a rubber break pad attached at the front that generates frictional force on the rear foot as the user walks. Optical encoders are used to measure the length of the user's step and any turning angle so that body position and orientation can be calculated. A mechanical tracker is used for monitoring the position and orientation of the head. These data are then used to correlate the visual display, provided by a HMD, with movement of both feet and head. A metal hoop positioned around the user's waist limits his forward or backward motion so that the user remains in the same place. A photograph of the device is given in Figure 119.

The complete system uses two PCs: one with a graphics accelerator for image generation, and a second for supervising the motion tracking of feet and head. To provide manual input to the VE, an input device such as the Haptic Master (see Section 6.2.2.5) can be mounted on the steel hoop. The walkthrough simulator can also be used in conjunction with a motion platform to support the simulation of movement across an uneven surface.

The researchers have used the walkthrough simulator in two experiments that examined the potential benefit of being able to simulate walking through a VE, as opposed to the current mode of flying via hand gestures. Specifically, these experiments looked at how walking and flying compare with respect to distance estimation (Iwata and Matsuda, 1992). In each case, three subjects were asked to follow a given path through a VE until they reached a marked goal. For the first experiment, six straight line paths were used and, after completing each, the subjects were asked to simply mark the position of the goal on a prepared worksheet. The second experiment used six paths in the shape of a four-sided figure and, after each completion, the subjects were asked to sketch the shape of the path so that the estimated distances could be measured. In both experiments, distances greater than ten feet were always overestimated, and lesser distances were frequently underestimated. However, the distance estimates given after walking a path were closer to the true distances than those given after flying the path.

Research has just started on improving a virtual staircase display. The original virtual staircase display used strings to pull a user's feet, but this approach led to problems in body stability. Researchers are now investigating the use of a 3 DOF motion platform to implement a virtual staircase display.

Future plans for the walkthrough simulator will include applying the simulator in the study of human behavior in emergency situations; in particular, it will be used in a refuge simulator for fire accidents.




Figure 119. Haptic Simulator

7.3 Passive Motion Interfaces

Passive motion interfaces will continue to play a large role in training simulators and entertainment applications, and visual displays and motion platforms will continue to be used as the primary means of supporting the illusion of movement. As stated earlier in this report, the traditional types of motion platforms used in simulators represent a mature and well-defined technology area that is not addressed in this report. Similarly, the traditional type of cabin simulator is not discussed in detail, save to mention that the application of VE interface technology to cabin simulators is relatively new, and the development of virtual cabin simulators will allow more general-purpose simulators to be built, and support the rapid prototyping of simulator designs. Working towards this end, Boeing Computer Services (see Section 6.2.3.1) and CGSD Corporation (see Section 6.2.3.2) are both investigating the development of a virtual cockpit instrument panel. Researchers at the Air Force Institute of Technology are developing a virtual cockpit using an HMD for display of both out-of-the-window images and the cockpit itself. In its Fusion Interfaces for Tactical Environments (FITE) Laboratory, Armstrong Laboratory is taking an augmented reality approach to the research and development of advanced pilot interfaces for the F-16. Here the pilot sits in a F-16 cockpit shell and is presented with both virtual and non-virtual visual and auditory displays. None of these effort incorporate actual motion cues, but rely on visual scenes to induce a sense of passive motion.

There is a new type of product on the market that offers a lower cost approach to providing passive motion cues. These products, typically based on the military's G-seats, limit the provision of kinesthetic cues to those that can be presented through a chair.

7.3.1 Commercial Products

The set of motion chairs currently on the market are predominantly intended for use in VE entertainment applications. Five such products are considered here. The characteristics of these products are summarized in Table 19, and then each is described in more detail below.

7.3.1.1 Cyber Air Base

From ViRtogo, Inc., the Cyber Air Base is a pneumatically powered 6 DOF motion chair. In its current form the motion is driven by either VHS, laser disc, or broadcast sources. The visual display is provided by Kaiser Electro-Optics Inc.'s VIM 500 HRpv and supported by head tracking. RS-485, RS-232, and four extra programmable analog and digital control pins allow for customization with special interface devices, such as a steering wheel or pedals. Special effects such as a "rumble and thump" generator are available as options. A photograph of the Cyber Air Base is given in Figure 120. The standard single



Photo courtesy of ViRtogo, Inc.

Figure 120. Cyber Air Base

Table 19. Characteristics of Commercially Available Passive-Motion Products


Product	Vendor	Source of Motion	Range of Motion	Additional Provided Equipment	Price
Cyber Air Base	ViRtogo, Inc.	Pneumatic reservoirs	26-29° pitch, 26-29° roll, 26-29° yaw	HMD, head tracker, optional rumble and thump generator	\$19,999
Cyberchair	RPI Entertainment	2-3 axis motion base	20-360° pitch, 30-360° roll (optional yaw movement)	Visual display (various), 3-D localized sound display, hand controllers, rumble and thump generator, air stream generator	\$14,000
CyberMotion Seat	CineMotion International plc	Pneumatic reservoirs	18°pitch, 18° roll, 125 mm heave	LCD flat panel display, sound display, hand controller, rumble and thump generators	\$55,000 (for 2-seat unit)
IntelliSeat	Torus Systems, Inc.	Pneumatic cushions	Not applicable	Controlled seatbelt, heated/cooled air streams, 12 scents	\$850
SIM245	Jessler Enterprises, Inc.	Brushless electric motors	45° pitch, 45° roll	HMD, hand controllers, rumble and thump generators	From \$19,000

seat is priced at \$19,999. A double seat version is also available with either a projector and screen or two VIM HMDs. A 4-6 seater version is being developed for release in the fourth quarter of 1996.

An interactive version of the Cyber Air Base is under development and expected to be released in the middle of 1996. The motion for this version of the chair is computer-controlled, and the user is provided with a programmable joystick for giving input commands.

7.3.1.2 Cyberchair

RPI Entertainment markets the Cyberchair. While primarily intended for the entertainment industry, this product is available to developers of other types of VE applications. The overall design of the chair is based on a flight seat mounted on a RPI Entertainment developed motion platform. For the standard model, the motion platform provides 20° pitch and 30° roll, though these can be increased to a full 360° as a special option. A three axis design is available to provide yaw movement. Rumble and thump effects are produced using bass speakers and vibration transducers. A variety of RPI Entertainment or OEM developed visual displays are available for use with the chair, including HMDs, projection, pull-up head-coupled, and lean-in head-coupled displays. The basic user input devices, mounted on a chair arm, are a joystick and a trackball. Various special effects can be provided, for example, air streams can be used to provide the sensation of a wind. A photograph of the Cyberchair is presented in Figure 121.



©1995 RPI All Rights Reserved.

Figure 121. Cyberchair

With respect to game support, over forty-four licensable games are available. Cyberchairs can be networked via a standard telephone interface to allow multiple users to share the same experiences. The price of a single Cyberchair is approximately \$14,000, depending on the configuration. For quantities of over three hundred, the price ranges from \$3,000 to \$62,000.

7.3.1.3 CyberMotion Interactive Motion Seat

The CyberMotion interactive attraction marketed by CineMotion plc includes two individual interactive seats incorporating a pneumatic motion control system that offers a simulation industry standard 3 DOFs. The motion system is self-contained, and located under the seat. It provides up to 18° of pitch and roll, and 125 mm of heave (lift).

Each CineMotion seat comes with a multi-function joystick coordinating both the motion system and the PC CD-ROM interactive game. The game is displayed and played in conjunction with a 12.1 inch LCD flat screen. Additionally, the motion system is connected with an audio control system that allows audio signals to automatically generate vibration and

conjunction with a 12.1 inch LCD flat screen. Additionally, the motion system is connected with an audio control system that allows audio signals to automatically generate vibration and shock effects. The system is controlled by a PC, and up to eight systems can be networked together. A photograph of the CyberMotion chair is given in Figure 122.

Marketed as a two seat module complete with individual coin operated mechanisms, and with the first CD-ROM interactive game included, the CyberMotion systems costs \$55,000. (The company makes no provision for providing the air compressor used in the pneumatic motion system, but will do so on request.)

CineMotion also markets the AirRide Passive Motion seat.

During 1997, CineMotion will be announcing further gameplay products in the sitting and standing positions, again using their internationally patented pneumatic motion system. Research developments also are underway for a low-cost motion seat for use in the home market.




Photo courtesy of CineMotion International plc

Figure 122. CyberMotion Interactive Motion Seat

7.3.1.4 IntelliSeat



Photo courtesy of Torus Systems, Inc.

Figure 123. IntelliSeat

The IntelliSeat, developed and marketed by Torus Systems, Inc., was designed as an alternative to hydraulic or electromagnetic motion systems. Instead, this system uses six individual pneumatic reservoirs placed in the seat and back cushions of the chair to provide the illusion of motion. The seat belt is controlled to give additional motion cues by tightening or loosening to indicate forward or reverse-direction accelerations. A panic button is provided that a user can push to immediately stop the chair's response to all computer-generated motion commands. A photograph of the IntelliSeat is provided in Figure 123.

The IntelliSeat system supports two non-motion sensory cues. The first of these, called TorusBreeze, is achieved by using micro fans to blow air streams that are strategically directed to "buffet" a user's ears and face. Small heating and cooling units can change the temperate of the air streams. The second non-motion sensory cue provides odors and is called Toruscent. Here a total of twelve scents, six in each chair arm, are stored in solid form and dispersed using the TorusBreeze. A proprietary activation system triggers the scent units to open and close so as to deliver a very faint essence chosen to enhance the VE imagery. An additional sensory

modality is planned for release in summer 1996. Called TorusMist, this subsystem will use a atomizer contained in the chair arm to spray small amounts of water on the user. The intent here is to make the user cool, rather than wet; providing a sensory cue for actions such as moving through bushes after a rainfall. The basic IntelliSeat is available for \$850.

Another version of the seat, called IntelliSeat340, is exclusively intended for use with Torus Systems Inc.'s Toruscope 360 special venue motion picture format system.

7.3.1.5 SIM245

Jesler Enterprises, Inc.'s SIM245 is another motion chair. This chair is powered by brushless electric motors, providing 45° pitch and 45° roll. Two user controls are mounted on the chair, these are a joystick with a dual fire switch and a turbo throttle with auxiliary button. The head mount provides support for the i-glasses, VFX1, CyberMaxx, VIM, or FS5 HMDs. Rumble and thump effects are provided by means of a 10 inch sub woofer and a sound amplifier. The system is controlled by a Pentium 90 MHz computer with a 16-bit sound card. A photograph of the SIM245 is shown in Figure 124.




Photo courtesy of Jesler Enterprises, Inc.

Figure 124. SIM245

The price for the SIM245 depends on the HMD selected and ranges, for a single unit, from \$19,000 to \$24,000. This includes one game, with additional games available at \$1,500 each.

7.3.2 Current Research and Development

Two efforts investigating the development and use of motion seats have been identified and are discussed below.

7.3.2.1 Denne Developments Limited

Under contract to the British Government, Denne Developments, Ltd. (DDL) is developing a G-seat, called the CyberSeat. G-seats are intended to overcome the limitation of motion platforms for supplying sustained acceleration forces by simulating motion through the application of pressure to the skin. As a low cost alternative to the traditional pneumatic actuators currently used in G-seats, DDL will use their patented PemRAM technology (see Section 7.3.2.1) to provide sustained forces through the cushions on the back and seat of a specially designed chair. Since this approach does not require large body movements to simulate user motion, it is expected to avoid the difficulties that heavy headsets can cause when used with conventional motion platforms.

A prototype CyberSeat is expected to be completed by late 1996, with a commercial product released in 1997/98.

7.3.2.2 Flogiston Corporation

Under a SBIR Phase II grant from NASA, Flogiston Corporation is developing a personal motion platform (PMP) that will be used for training astronauts in EVA procedures. The PMP will provide visual and auditory interfaces, in addition to a motion interface to a VE. The major element of the PMP is the Flogiston Chair, a reclining device developed under previous NASA funding that supports the body in a neutral posture. This chair is mounted on one of two motion platforms: a 3 DOF platform that uses electromagnetic ram actuators, or a 6 DOF platform driven by pneumatic actuators. The visual interface, attached to the chair, will employ a device that provides a high resolution image and supports peripheral vision. The auditory system will consist of off-the-ear headphones, providing 3-D localization of sound, and moving mass vibrators positioned on the chair and the motion base. Together, these two types of auditory interfaces will be capable of providing frequency signals ranging from 0 to 20 kHz. The user controls movement through the VE by means of 6 DOF pucks, mounted on the chair to lie under each hand. A photograph of the current PMP is provided in Figure 125.

Supporting software will allow integration of the PMP with VE systems on Unix and NT Windows machine environments. It will include a generic behavioral model capable of interpreting user input and mapping simulation-generated motion, auditory, and visual cues to the appropriate stimuli.

The current work is centered around refining and extending an initial PMP developed under an earlier SBIR Phase I grant. This work entails developing new software for the motion platform, in particular, developing a new motion control algorithm and drivers for different motion platforms. Tests of the PMP are being conducted that include refining PMP motion cuing based on the stress levels induced by particular motion platform movements. Future work is expected to include the development of a haptic feedback device to replace the current 6 DOF hand controller.

In addition to delivering the PMP to NASA, Flogiston will market a commercial version of the system called the Flostation. This product is expected to come to market in late 1996.

7.4 Summary and Expectations

Flying by means of some hand-held control device will always be a common method of representing movement through a VE. It is a low-cost approach, one that is not physically taxing on the user, and suitable for many types of applications. While some gyroscope devices are commercially available and are being used in entertainment applications, studies on the effects of frequent or prolonged use of these devices on the user are needed. Nonetheless, the




Photo courtesy of Flogiston Corporation

Figure 125. Personal Motion Platform (PMP)

use of gyroscopes, hang gliders, and probably similar devices yet to appear, also is expected to continue. In addition to entertainment applications, such motion interfaces are likely to see limited use in specialized types of applications, such as those dealing with space or underwater environments. The potential role of interactive motion platforms has not been fully investigated as yet and, again, there are human factors and safety issues to be addressed. It is believed, however, that this type of responsive motion platform is going to see wide use, in both traditional motion platform application interfaces and in innovative ways.

To date, research into interfaces that allow some type of active user locomotion, albeit restricted, has seen little progress. Six prototype interfaces systems have been identified, three developed by Sarcos Research Corporation, and one each by the Institute for Simulation and Training, University College London (UK), and the University of Tsukuba (Japan). Each group of researchers has taken a different approach using, for example, a unicycle, treadmill, walking in place, and modified roller skates. In addition, STRICOM is sponsoring the design of three new prototypes, one of which is expected to be selected for further development. One of these designs employs a 3-D treadmill and the other two are based on the concept of movable foot plates. It is worth pointing out that all the current approaches to active locomotion, except one, are mechanical in nature. While each approach limits user movements in varying degrees, there are inherent differences in the way the mechanical and non-mechanical approaches limit and facilitate user movement. For example, technology advances in tracking and recognizing body movements, and in tracking range, will allow the current non-mechanical moving in place approach to be expanded to allow a wide range of user movement, but by itself will never support conditions of varying surface conditions or obstacles. Mechanical approaches, on the other hand, limit the user to a restricted set of movements and those movements may only approximate real motions. However, some mechanical approaches are capable of simulating a variety of surfaces, including such obstacles as stairs.

It is expected that research and development on locomotion interfaces will continue. Probably this work will be funded largely by the DoD to meet particular training requirements, but it should result in spin-offs applicable for non-DoD applications. While some advanced prototypes may see trial use in specific applications during the next five years, such systems are not expected to come into practical use within this timeframe. As advanced prototypes are developed, researchers will need to investigate many issues, such as the importance of fidelity to normal human motion for particular applications, and human factors concerns.

As discussed in the introduction to this section, traditional cabin simulators are a well developed field that has been excluded from consideration. As a result, the material presented for passive motion interfaces focuses exclusively on motion chairs: five commercial products and two being developed in research efforts. The majority of these devices provide inertial displays in which body mass is moved, using electromagnetic or pneumatic actuators to provide motion in 3 to 6 DOFs. The two motion chairs that use a non-inertial approach both provide pressure through the seat and back chair cushions to induce the sensation of motion (the IntelliSeat uses pneumatic actuators and the CyberSeat uses PemRAM electromagnetic actuators). With one exception, the Personal Motion Platform, all current motion chairs are primarily

intended for entertainment applications and most of the commercial systems provide special effects such as breezes blowing across a user and scents. There are no major technical challenges in this area. Further products are likely to come to market in the next few years and the only significant issue is one of matching device capabilities, and price, to application needs.

8. OLFACTORY INTERFACES

The olfactory interface is one of the least developed areas within the field of human-computer interaction. There are a number of reasons why this has been so, the main reasons being the lack of useful applications and the current societal mores associated with olfaction. However, with the advent of new VE technology, olfactory interfaces are now seen as a valuable sensory cue for applications such as fire-fighting and surgical training.

While the input or sensing device for an olfactory interface is not solely within the domain of VE technologies, it is a necessary component for the development of VE olfactory systems. These devices are commonly referred to as artificial or electronic noses and are used to collect and interpret odors. There are three basic approaches to sensing technology: gas chromatography, mass spectrometry, and the use of chemical sensor arrays. These are used in a range of applications, such as chemical and biological warfare detectors (used in the Gulf War), and product quality control. It is likely that the same types of technology are suitable for use in acquiring data on odors to be used in a VE. The focus of this section, however, is on systems that can deliver olfactory cues in a VE.

Odorant storage is, perhaps, the most mature of the various technologies required for an olfactory delivery system. Odorants can be stored in a number of ways, including as liquids, gels, or waxy solids. The most popular storage method for previous and current VE-related work seems to be to microencapsulate odorants. This method is the basis of scratch-and-sniff patches. Droplets of liquid (ranging in size from 10-1,000 μm) are encapsulated in a wall of gelatin. They can be printed using silk screen techniques, allowing multiple odors to be printed onto a flat surface. Typically, the odorant is released by subjecting the particle to mechanical shear, or melting the gelatin wall. Microencapsulation offers the advantages of discrete metering of odorant dosage, stability at room temperatures, and the unlikelihood of messy spills. Released odors must then be presented to the user. At present, the major methods include air dilution olfactometry, breathable membranes coated with a liquid odor, and a system of liquid injection into an electrostatic field with air flow control. A summary of the strengths and weaknesses of the various delivery technologies is given in Table 20.

Olfactory delivery systems for VEs, however, require more than odor storage and display. They also need to clean the air input, select odorants for display, and evacuate and clean exhaled air. The greatest obstacle to this is in controlling the breathing space for the individual; for example, it is necessary to accurately control odor intensities, quickly flush

Table 20. Olfactory Delivery Technologies^a

Storage Technologies	Presentation Technologies	Advantages	Disadvantages
Liquid	- Unpowered evaporation: Saturated cotton balls Breathable membranes Permeation tubes Bubble chambers	- No power - Inexpensive	- Bulky - Odorants clumsy to handle
	- Heat induced evaporation	- Inexpensive	- Power hungry
Gels	- Electrostatic evaporation	- Good for large spaces - Materials easier to handle	- Never miniaturized - Requires higher voltages
Microencapsulation	- Mechanical release	- Could be valveless - Materials easy to handle	- Mass production technology - Impractical for small lots
	- Heat release	- Could be valveless - Materials easy to handle	- Mass production technology - Impractical for small lots
	- Valve design options: No valves Off-the-shelf valves	- Smaller, cheaper	- Intercontamination of odors
		- Mass produced	- Bulky, power hungry - Fast or precise, not both
	Ink jet printer nozzles Microvalves	- Precise control	- Single units large because of packaging
		- Potentially fast & small	- Must make custom minifolds to get greatest miniaturization

a. Based on Krueger (1995, 1996).

an odor from the breathing space when a particular odor cue is no longer required, and prevent any contamination by persistent odors. Krueger (1995) identifies several ways of presenting odors that attack this problem with varying degrees of encumbrance to the user:

1. A sealed room with a precise air filtration system.
2. An unsealed cubicle that directs treated air toward the user's face and that provides a collection vent behind his head to evacuate the odorized air he exhales. This still requires some general air filtration system for the room housing these cubicles.
3. A completely sealed pod in which only treated air is breathed and exhaled air is continually evacuated.
4. A tethered mask that can be used in a general purpose room by either a seated or stationary standing user.
5. An untethered system that would consist of a belt pack and tubes running to and from a mask in a HMD.
6. An untethered system that is completely incorporated into the HMD itself.

In addition to differing in the degree in which they encumber the user, these alternative ways of presenting odors differ greatly in such factors as cost, space, and support requirements.

Why odors? There is evidence that odors can be used to manipulate mood, increase vigilance, decrease stress, and improve retention and recall of learned material. One recent experiment demonstrated that a peppermint odor gave superior performance to a lavender odor or no odor at all in spatial visualization and perception tasks (Krueger, 1995). Knasko and Gilbert (1990) found that even the suggestion of odors described as pleasant, unpleasant, or neutral can lead subjects to give self-reports of pleasure and induce a more positive mood. In this experiment, the number of reported physical health symptoms differed as a function of the hedonic quality of the feigned odor; the condition with the feigned pleasant odor reported the fewest number of physical symptoms. Although subjects in the unpleasant odor condition predicted higher task performance, actual performance did not differ across the conditions. Also, as with auditory cues, it is possible that odors can be used for sensory substitution, representing phenomena that have no smell or purely abstract information.

This section starts with a brief overview of the human olfactory sense, followed by descriptions of two commercial products. The discussion then moves on to review research efforts in this area. As usual, the section closes with a summation of likely developments in the near future.

8.1 The Human Olfactory Sense

When a human sniffs an odor, molecules carrying the scent are captured by the receptor neurons in the nasal passages. The cells that become excited fire pulses that travel through axons to a part of the cortex known as the olfactory bulb. The number of activated receptors indicates the intensity of the stimulus and their location in the nasal passage conveys the nature of the scent. Each scent is identified by a pattern of receptor activity, which in turn is transmitted to the bulb.

The bulb analyzes each of the input patterns and then synthesizes its own message, which it transmits to the olfactory cortex. These new signals are sent to many parts of the brain where they are combined with signals from other sensory systems. The result is a contextual perception of the odor that is unique to each individual. This is, however, an incomplete account of olfaction. There are a number of questions that remain unanswered. For example, how does the brain distinguish one scent from all the others that accompany it and how does the brain generate a pattern when some receptor signals are missing? The University of California at Berkeley, Harvard University, and Yale University are all researching the underlying basic science issues in mapping out the human olfactory system. The University of California is focusing on mapping the spatial patterns that the brain recognizes as smell, while the Harvard and Yale research has centered on mapping the DNA of the receptor sites.

As reported by Krueger (1995), there are two senses that are closely related, but distinct, from olfaction. One is taste. The second relates to the tactile sensors in the nose (and also in the mouth and throat) that detect hot and cold, irritation and pain sensations. The sharp smell of ammonia, for example, is actually a tactile sensation that is reported to the brain through the trigeminal nerve (the fifth cranial nerve) rather than through the olfactory nerve (the first cranial nerve). In general, the greater the trigeminal component in an odor, the faster it is recognized, although the perception of oral heat does have a long lag.

Other sensory systems play an important role in human olfaction. As reported by Zellner, Bartoli, and Eckard (1991), for example, humans may correctly identify only one third of odors in the absence of input from other sensory systems, such as vision. These researchers review their own and others' work in assessing the role that color cues have on odor identification. Overall, the findings show that when the appropriate color cue is presented with an odor, both the accuracy and speed of identification improve. Conversely, an inappropriate match of color cue can lead to reduced accuracy and longer response time.

The human capability to detect odors is quite sensitive, capable of detecting odorants in concentrations of one part per million, or even one part per billion, depending on the odor in question. Data on identification thresholds and reaction times for a range of different odors is given in several sources, for example (Overbosch et al, 1989), (Naus, 1985), and (Laing, 1986). Increases in concentration are far more likely to be detected than decreases. Krueger (1995) reports that the smallest detectable change is a 15% to 30% increase in concentration; perceived magnitude is not linear with changes in concentration,

but closer to a logarithmic relationship. Further, Krueger makes the point that many studies have shown that humans can only reliably identify such gross measures as: barely detectable but not identifiable, barely identifiable, clearly present, strong, and very strong. It is known that response of smell receptors is time, temperature, and humidity dependent. But many other factors play a role. Segal et al (1995) report that there appears to be a genetic influence on odor identification for males, but not for females. These researchers also found that there is a curvilinear age trend for males, but, again, not for females. Also, the acuity of the sense of smell is subject to change. This change can arise due to physiologic or pathologic reasons. In most cases, however, prolonged or repeated exposure to an odor can result in adaptation that reduces detection.

8.2 Commercial Products

In addition to an entertainment-oriented motion chair that releases odors into an uncontrolled air space (see Section 7.3.1.4), only two commercial olfactory delivery systems intended for use in VEs have been identified.

8.2.1 BOC Group Olfactory Delivery System

The BOC Group plc, in the UK, market an olfactory delivery system to organizations such as VE entertainment and video game producers. Their patented approach is based on dissolving odorants in an environmentally friendly, high pressure solvent, such as carbon dioxide, and then delivering the resultant gas via an air stream blown at the user. The actual delivery system is computer-controlled and delivers dose levels down to the parts per billion level. It can be attached to a HMD. BOC Group plc works with various fragrance houses to enable them to deliver a very wide range of odors.

8.2.2 Smell-Enhanced Experience System

Ferris Productions, Inc. developed the first commercial VE-related olfactory delivery system, integrated in an entertainment-based system called the Experience System. The Experience System includes a NASA-developed zero gravity position chair, 3-D spatial sound, 3-D visuals delivered by a HMD, and the olfactory capability.

The olfactory system stores up to seven odors in liquid form in separate canisters. Odors are generated by releasing controlled amounts of an odor into an air stream produced by a 20 psi air compressor. The scented air stream then is delivered to the user via a small hose pointed towards his nose. The system can be used with an uncontrolled air space, or the user can wear a mask that can be integrated into any HMD. The odors introduced into the air space are expected to clear within about one quarter of a second. In addition to its use with the Experience System, the olfactory delivery system is available as an independent unit. It is controlled by a stand-alone, microchip-based system that not only turns a selected odor on or off, but controls the strength of a generated odor. The price of the complete Experience System is \$11,999. The price for the stand-alone olfactory delivery system starts at \$4,000.

8.3 Current Research and Development in Olfactory Interfaces

The earliest known work in providing olfactory input for VEs was an internal research study performed at the Southwest Research Institute in 1993. This work led to the development of a prototype odor producing hardware system called DIVEpak. Controlled by a microcomputer, this system could deliver eight different odors. The (essential) oil-based odors were encapsulated and contained in a cartridge. When released, the capsules were ruptured using heated motors and then air was blown across the liquid odorant to let the odors evaporate into the air stream. Trials with the prototype were partially successful and design modifications were defined to resolved problems found with the DIVEpak. At the completion of the study, further work in the area was placed on hold pending active market interest.

Four groups have been identified as currently pursuing research in the area of olfactory delivery systems. The E. Piaggio Bio-Robotic Laboratory at the University of Pisa is developing a VE with integrated olfaction for telemedicine applications. This work includes the development of an odorant capture device (called a smell camera) to record odor patterns for regeneration and an olfactory delivery system for the odor regeneration. Dr. Clifford Bragdon at the National Aviation and Transportation Center, Dowling College, is developing a so-called multimodal simulation system that will support a variety of transportation modalities, olfactory stimuli, and 3-D sight and sound. Further details on these three efforts are not available. The remainder of this subsection discusses the work of the remaining two groups of researchers, those at Artificial Reality Corporation and at Marketing Aromatics, Ltd.

8.3.1 Artificial Reality Corporation

Sponsored by ARPA, the Artificial Reality Corporation (ARC) is conducting a feasibility study for the inclusion of olfactory interfaces in VEs. Part of the plan for this work is to review the state-of-the-art in olfactory sensing and odor delivery to individuals, and to assess the basic science, technology, techniques, and products that are available on the market. This part of the work has been completed, see (Krueger, 1995). Additional work includes a series of studies aimed at ascertaining the effects of odors on the acquisition of skills related to surgery and addressing such questions as: Do appropriate olfactory stimuli add to a sense of presence in a VE? and, Do appropriate odors improve efficacy of VE training? Experiments are now looking at the impact of olfactory stimulation on the acquisition of fine motor skills. Additionally, the researchers are negotiating with the developers of some surgical simulators to the possibility of developing an integrated system to support further experimentation.

Integration of an olfactory capability with a surgical simulator requires a number of specialized odorants that have not previously been available. Here ARC is working with two other companies, Monell Chemical Senses Center, and International Flavors and Fragrances, to develop the necessary odorants in liquid form. Odorants for human body, blood,

and liver odors have already been developed, but five or six more odors are deemed necessary to represent the common odors experienced in surgery. A prototype of this first olfactory system using only commercially available clean odors and the three current special odorants could be ready by Fall '96. This prototype is expected to be an environmental unit such as a pod or booth where odors are introduced through the floor and vented through the ceiling using a closed air system. This type of controlled breathing space was chosen because it is known to work, although it can be very expensive. Subsequently, the researchers will investigate the capabilities of odor delivery systems that are less intrusive on the instructional environment. The researchers also are considering modifying a CAVE system to include an olfactory system.

ARC has identified a design option that may allow miniaturizing a delivery system, enough that it may fit inside an HMD. Such a portable olfactory system requires miniaturized and lightweight components with low power requirements, and ARC is currently examining candidate technologies. For example, ink-jet printer nozzles are being considered for odorant delivery since these will allow precise control of some odorants. Memory metal valves and electrostatic diffusion delivery technology also are under consideration. When a prototype has been developed, probably by mid 1997, it will be used to study olfactory perception in the context of physical behavior and to develop a testbed for medic training.

Potential future work will address the use of olfactory stimulation for telepresence medical applications. Here odors will be measured at one site and electronically transmitted for reproduction at the surgeon's remote site. Chemical sensors are not yet fast enough to detect rapid changes in odors and the researchers plan to look at the use of a mass spectrometer that operates continuously for the measurement element of this work. The effort will include identifying which odors are relevant to medical applications and pick a set of these for demonstration. In addition, odors at surgical procedures will be recorded for use with video tape presentations.

8.3.2 Marketing Aromatics, Ltd.

Marketing Aromatics, Ltd. is working on a technology for olfactory delivery systems that is intended to meet three critical physical criteria: a rapid rise/decay of olfactory stimuli, provision of a wide palette of odors, and microprocessor control of the delivery process. The technology itself is a spin-off from other company work and employs aromatic oils that are effectively vaporized to an almost molecular level, thus allowing precise control of minute amounts of vapor. The conversion is very rapid, with vapor generation occurring in the order of milliseconds. The vapor can be delivered, via an air stream, to the user in a number of ways, for example, by applying an electric charge and then directing the vapor using an ionic wind. The user's air space is controlled, using a mask that can be rapidly evacuated.

Little information is currently available about the actual technology. The key part of the technology, however, is the patented vaporization procedure. The initial delivery system is expected to include around twenty odors, provided by nozzles mounted inside the mask close to the user's nose. The actual choice of delivery and evacuation systems has yet to be made, but the researchers anticipate an odor decay time of less than 1 second. The technology is expected to become commercially available within the next two to three years. Marketing Aromatics, Ltd. see their major consumers being organizations that develop VE entertainment applications and fragrance houses.

In other work, Marketing Aromatics, Ltd. are looking at the use of their olfactory generation technology for large air spaces, such as shops, offices, and airports.

8.4 Summary and Expectations

It seems likely that some olfactory delivery systems for VEs will be developed in the next few years, but these are expected to be largely prototype systems intended for research and experimental purposes. Problems to be solved include the mechanical ones associated with odor storage, selection, regeneration, and breathing space control. Early devices are likely to be too large and heavy for prolonged use, especially if air tanks are required to provide a fresh air supply. Another impediment is the scarcity of suitable odorants: the types of odors likely to be required for use with VEs are unlikely to exist in the standard repertoires of fragrance companies and will take time, and funding, to develop. For these reasons, and others, it is doubtful that a practical olfactory delivery system will be derived from existing technology with the next five years.

9. CONCLUSIONS

At the present time, visual, tracking, and primary user input interfaces are the ones best suited for practical VE applications. In each of these cases, there is a solid basis of commercial products for potential users to choose from. Auditory and haptic interface technologies currently are largely restricted to research applications, but are on the verge of becoming ready for use in practical applications where such interaction is deemed essential. Although widescale usage of auditory interfaces is expected to precede that of haptic interfaces, it is still some time away for both technologies. With respect to full-body motion interfaces, there are several entertainment systems that support limited types of highly specialized movement. Support for more general types of movement still is exclusively a research topic with a variety of different approaches being investigated and motion interfaces systems are unlikely to become suitable for practical use within the next three to five years. Current work on interfaces for passive motion is focusing on a new breed of motion chairs, largely intended for the entertainment market. Olfactory interface technology is the least mature of all the technologies discussed here and another unlikely to see practical usage within the three to five year timeframe.

All current VE interface technologies suffer from some limitations, even the more mature visual, tracking, and primary user input technologies. In no instance does the interface technology match human capabilities for the relevant sensory modality.

In the case of visual interfaces, HMDs and CAVES (typically using projection screens and passive glasses) are the only means of achieving an encompassing visual volume. HMDs, much more widely used than CAVES, suffer from several problems, with the most serious limitations being:

- Inadequate display update rates when responding to user head movements.
- Inability to provide both high resolution and a broad field of view.
- Weight that imposes an inertial burden and low levels of comfort that prevent prolonged use.

All these problems are well recognized and the first two are likely to be substantially reduced in the next few years through advances in LCD technologies. While smaller, lighter weight displays will help to reduce overall HMD weight, the necessity for bulky optics means that weight will continue to be a problem. A former problem, the expense of commercial HMDs, is becoming less serious as more low cost devices are becoming available,

although these require the user to make some compromises in resolution and/or field of view.

So far passive glasses have not been widely used in VE applications, although new microelectronic fabrication techniques for creating polarizing filters at the pixel level may change this trend. Shutter glasses are quite widely used, usually with CRTs or projection displays. Here again, advances in LCD technology are likely to see an impact as LCD displays with faster switching time will help in reducing crosstalk problems. There is much research and development in the area of autostereoscopic displays and a small number of products is likely to come to market in the next two to three years. While these displays offer the advantage of not requiring any encumbering head gear, glasses, or head tracking devices, they also have some current limitations. The primary limitation, that users are restricted to a limited viewing area, is likely to be reduced with the development of flat panel displays with higher resolution and the simultaneous display of larger numbers of perspective views. Retinal displays are a new topic of research and development. While they have the potential for providing a fully encompassing visual display without the weight and limited resolution and field of view of current HMDs, it will likely be some years before black-and-white retinal displays come to market, and longer for color displays.

Systems for tracking head, hand, and body movements are available and many have seen widespread use. Even so, low latency, high accuracy systems for tracking in noisy, unprepared environments do not exist. The most serious shortcoming of current technology is the following:

- Inherent limitations in some combination of accuracy, intrinsic latencies, working volume, susceptibility to interference or obscuration, and cost.

Again, these are well-recognized problems that are expected to be the focus of near-term research and progress, especially for magnetic trackers, is expected. The most significant improvements in tracking performance, however, are expected to come from the use of hybrid trackers where many of the limitations inherent in a particular technology can be overcome. Research in the development of such hybrid trackers is underway. Although there are no products commercially available as yet, these are expected to start appearing within the next couple of years. Wide-area trackers are another area where commercial products are unavailable and, with only limited research being performed, this type of tracking interface is not expected to see widespread use any time soon.

Eye tracking also is a less mature type of tracking technology, largely because traditionally it has had a limited range of applications and, therefore, has attracted little research interest. In this case, the major problems appear to relate to:

- Limited accuracy and intolerance to user head movements.

The increased use of multimodal interfaces (in both VE and non-VE applications) that can benefit from the ability to monitor the direction of the user's gaze, however, is opening up

new potential markets that should encourage further development of this type of interface technology.

A number of 3-D sound processors that can be used in VEs are commercially available. These range in capability from systems available for use with PCs, to high-end professional audio systems. However, a number of questions need to be answered and further research done before virtual audio can become a practical tool. Serious limitations are the following:

- Inability to represent sounds as being located in front of the user and to adjust sound spatialization to head movements.
- Inadequacies in acoustic signal generation.

Near term work is expected to focus on these areas, continuing to improve the realism and full-surround capabilities of the technology. Crucial support for this work will come from the development of improved algorithms, based on a more thorough understanding of how humans perceive sounds. As digital signal processing becomes less expensive, virtual audio is likely to become more widespread. This is already happening to some extent with many dedicated game systems, major computer companies, and audio chip manufacturers licensing low-end virtual audio technology. As a result of increasing availability and the lower cost of technology, these types of interface are expected to become a common component of VE systems within the next five years.

The development of glove-based devices for user input is an area of current growth. All the interface systems on the market are relatively new products and at least two additional products are expected to become available by mid 1996. The current set of products do allow the use of natural hand gestures for certain, limited interactions with a VE but the primary shortcoming remains:

- Limited joint resolution and poor discrimination between gestures.

While improvements in sensor technology will help to reduce this problem, it is likely that advances in software-based gesture recognition will play a more important role. Gloves already are a fairly common VE input device but their use is expected to become more widespread as gesture recognition capabilities improve. There seems to be little ongoing research looking at the use of exoskeleton-based devices and these are not expected to be widely used, but limited to highly specialized applications.

A fairly diverse range of 3-D pointing devices is available. These products represent mature technology and, while new products may appear over time, no major changes in this area are expected.

Tactile and force feedback interfaces for VEs have been able to exploit previous work in the areas of, respectively, sensory substitution devices for the disabled and teleoperation. Both represent active areas of research and development. In the case of tactile interfaces, researchers are investigating how to provide contact force, slip, texture, vibration,

and thermal sensations. Products intended to simulate contact forces that occur when a user touches a virtual object and that provide temperature feedback are already commercially available. The ability to support other types of tactile sensation is more problematic. Although prototype devices exist, each tends to be specialized to one particular type of sensation. Moreover, all existing devices, both commercial products and prototypes, limit the presentation of sensation to a small area, usually the fingertip, and are unlikely to be able to scale up. While these devices are relatively small and lightweight, at least compared to HMDs, they are encumbering to some extent and can constrain finger movement. In addition to shortcomings in tactile interface hardware, much work is still needed in developing the software models needed to drive the generation of tactile signals. Consequently, the major limitations in the area of tactile feedback can be summarized as follows:

- Limitations in the ability to represent surface characteristics such as texture, local shape, and slip.
- Inability of devices to present a range of tactile sensations.
- Limitation of tactile feedback to small areas.
- Lack of models and algorithms for efficient generation of tactile signals.

As stated, this is an active area of research and much progress is expected over the next few years. Nevertheless, although several prototype applications are expected, tactile interfaces are unlikely to see common use within the next two to three years.

The majority of current force feedback devices can be distinguished as exoskeleton devices that deliver forces to the shoulder, arm, or hand; tool-based devices that deliver forces to the hand via a knob, joystick, or pen-like object held by the user; thimble-based devices that deliver forces to the user's fingertips; or robotic graphics systems that move real objects into place to provide natural forces to the user. Here again, each type of device is limited in the type of interactions it can support, in this case largely because of the intrusive nature of each device. Consequently, although several devices are on the market, each provides very different capabilities and is suitable for different types of application and, as yet, these devices have only seen limited use. The serious limitations of force feedback interfaces are, in many respects, similar to those given for tactile interfaces:

- Inability to provide force feedback for a variety of different VE interactions.
- Limitation of force feedback to a restricted number of joints.
- Intrusive nature of force feedback devices and their constraints on user movement.
- Lack of common models and algorithms for efficient generation of kinesthetic signals.

Burdea and Zhuang (1991) cite deficiencies for teleoperator systems that also apply here. These include the inadequacy of current actuators, coupling between degrees of freedom, and high system complexity from mechanical design, hardware, and control software issues. Since this is an active area of research, considerable technology advances are expected to occur in the next three to five

years. In the interim, force feedback interfaces are unlikely to see much practical use, with the possible exception of those used in medical applications.

A number of approaches and devices have been developed to facilitate a user "moving" through a VE. The simplest, and most common of these, is for the user to point in the desired direction and for the visual scenes to be adjusted accordingly. A number of entertainment systems provide highly specialized interface devices allowing, for example, the user to simulate hang gliding or sledding. Unfortunately, there has been little progress in providing interfaces that allow a user to simply walk or run through a VE. Of course, active self-motion within a small area (for example, 10 x 10 feet), over a uniform surface that provides the necessary haptic cues presents no problem. Technology that can support a user moving through a large area or across a surface with varying characteristics, however, has only recently begun to be investigated. A number of diverse designs for interface systems have been proposed and a few prototypes built, using both mechanical and non-mechanical approaches. While such systems may see use as advanced prototypes, none are expected to come into common practical use within the next three to five years. The potentially large entertainment market also has fostered the development of passive motion interfaces. In the last year, several motion chairs have been developed that employ techniques ranging from inflatable chair cushions to motion bases to provide the user with a sense of motion. These devices may become widely used for a diverse range of low-cost simulators.

Three commercial systems that support the use of olfactory cues in entertainment applications are available, two of these providing controlled release of odors. There are no olfactory delivery products that support a controlled air space. A small number of research efforts are underway and at least two prototype systems are being developed. While more prototype systems might be developed in the next few years, this technology is not expected to become practically available in the near future. Some problems that are being addressed by ongoing research include:

- The encumbering nature of delivery systems and the need for miniaturization of systems components.
- Difficulty in controlling the user's breathing space.

While the ability to provide olfactory cues may be important for specialized applications, such as surgical training, the utility of such cues remains to be demonstrated.

In addition to further research and development on actual interface hardware and software, all the areas of interface technology discussed in this report will benefit from a better understanding of the role of sensory cues and human perceptual issues. This improved understanding not only is required to know how sensory cues can be delivered or simulated, but when and how they should be used. This is not to say that full fidelity of sensory cues is the ultimate goal. Even if achievable, high levels of fidelity would be expensive and not always desirable. What is needed is to determine the fidelity required for specific applications and how best to satisfy those requirements.

One issue that seems to have been ignored so far is that of usability. VE interface technology is primarily concerned with human-computer interaction and yet there have been no reported evaluations of the usability of particular VE interfaces. This type of study is keenly needed to guide both the use of existing interface systems and the ongoing development of new systems.

REFERENCES

- Adachi, Y., T. Kumano, and K. Ogino. 1995. "Intermediate Representation for Stiff Virtual Objects." In *Proc. VRAIS '95, IEEE Virtual Reality Annual International Symposium*, Research Triangle Park, NC, March 11-15, pp. 203-210.
- Adelstein, B.D., E.R. Johnson, and S.R. Ellis. 1995. "Dynamic Response of Electromagnetic Spatial Displacement Trackers." To be published in *Presence*.
- Akamatsu, M. 1994. "Touch with a Mouse. A Mouse Type Interface Device with Tactile and Force Display." In *Proc. 3rd IEEE International Workshop on Robot and Human Communication*, Nagoya, Japan, July 18-20, pp. 140-144.
- American National Standards Institute. 1969. *American National Standard Specification for Audiometers S3.6*. New York: ANSI.
- Applewhite, H. 1994. "A New Ultrasonic Positioning Principle Yielding Pseudo-Absolute Location." In *Virtual Reality Software and Technology: Proceedings of the VRST '94 Conference*, Singapore, August 23-26, pp. 175-183.
- Azuma, R. August 1994. *Registration Errors in Augmented Reality*. NSF/ARPA Science and Technology Center for Computer Graphics and Scientific Visualization. Available on the WWW at URL <http://www.cs.unc.edu/~azuma>.
- Barfield, W. and T.A. Furness, III. 1995. *Virtual Environments and Advanced Interface Design*. Oxford University Press: New York, NY.
- Begault, D.R. 1995. "Virtual Acoustic Displays for Teleconferencing: Intelligibility Advantage for "Telephone Grade" Audio." In *Proc. 98th Convention of the Audio Engineering Society*, New York.
- Begault, D.R. 1994. *3-D Sound for Virtual Reality and Multimedia*. AP Professional: Boston, MA.
- Begault, D.R. 1993. "Head-Up Auditory Displays for Traffic Collision Avoidance System Advisories: A Preliminary Investigation." *Human Factors*, 35(4), pp. 707-717.

- Begault, D.R. and T. Erbe. 1993. "Multi-Channel Spatial Auditory Display for Speech Communication." In *Proc. 95th Convention of the Audio Engineering Society*, New York.
- Begault, D.R. and E.M. Wenzel. 1993. "Headphone Localization of Speech." *Human Factors*, 35(2), pp. 361-376.
- Benson, A.J. 1990. "Sensory Functions and Limitations of the Vestibular Systems." In R. Warren and A.H. Wertheim (Eds). *Perception and Control of Self Motion*. Lawrence Erlbaum Assoc.: Hillsdale, NJ.
- Berger, E.H. 1981. "Re-Examination of the Low-Frequency (50-1000 Hz) Normal Threshold of Hearing in Free and Diffuse Sound Fields." *Journal of the Acoustical Society of America*, 70, pp. 1635-1645.
- Bles, W. and J.M.B. Vianney de Jong. 1982. "Cervico-Vestibular and Visuo-Vestibular Interaction." *Acta Otolaryngol*, 94, pp. 61-72.
- Boff, K.R., L. Kaufman, and J.P. Thomas (Eds). 1986. *Handbook of Perception and Human Performance*. John Wiley and Sons: New York, NY.
- Bolanowski, S.J., G.A. Gescheider, R.T. Verrillo, and C.M. Checkosky. 1988. "Four Channels Mediate the Mechanical Aspects of Touch." *Journal of the Acoustical Society of America*, 84(5), pp. 1680-1694.
- Brandt, T., J. Dichgans, and E. Koenig. 1973. "Differential Effects of Central and Peripheral Vision for Egocentric and Exocentric Motion Perception." *Experimental Brain Research*, 16, pp. 476-491.
- Brooks, T.L. 1990. "Telerobotic Response Requirements." In *Proc. International Conference on Systems, Man, and Cybernetics*, Los Angeles, CA, November 4-7, pp. 113-120.
- Brooks, F.P. Jr., M. Ouh-Young, J.J. Batter, and P.J. Kilpatrick. 1990. "Project GROPE—Haptic Displays for Scientific Visualization." *Computer Graphics*, 24(4), pp. 177-185.
- Bryson, S. 1992. "Measurement and Calibration of Static Distortion of Position Data from 3D Trackers." In *Proc. SPIE Vol. 1669, Stereoscopic Displays and Applications III*, pp. 244-255.

- Bryson, S. 1993. "Effects of Lag and Frame Rate on Various Tracking Tasks." In *Proc. SPIE Vol. 1915, Stereoscopic Displays and Applications III*, pp. 155-166.
- Burdea, G., R. Goratowski, and N. Langrana. 1995. "Tactile and Force Sensing for Computerized Hand Diagnosis." In K. Morgan et al (Eds), *Interactive Technology and the New Paradigm for Healthcare: Proc. Medicine Meets Virtual Reality III*, San Diego, CA, January 19-22, pp. 60-69.
- Burdea, G. and P. Coiffet. 1994. *Virtual Reality Technology*. John Wiley and Sons: New York, NY.
- Burdea, G. and J. Zhuang. 1991. "Dextrous Telerobotics with Force Feedback—An Overview. Part 2: Control and Implementation." *Robotica*, 9, pp. 291-298.
- Cholewiak, R.W. and C.E. Sherrick. 1981. "A Computer-Controlled Matrix System for Presentation to the Skin of Complex Spatiotemporal Patterns." *Behavior Research Methods and Instrumentation*, 13(5), pp. 667-673.
- Clymer, M. and G. Graves. 1994. "New Approaches in Magnetic Sensing for Tracking Devices." In *Proc. SPIE Vol. 2177*, pp. 288-296.
- Corso, J.F. 1963. "Age and Sex Differences in Pure Tone Thresholds." *Archives of Otolaryngology*, 77, pp. 385-405.
- Cutkosky, M.R. and R.D. Howe. 1990. "Human Grasp Choice and Robotic Grasp Analysis." In *Dextrous Robot Hands*, pp. 5-31. Springer-Verlag: New York, NY.
- Davson, H. 1989. *Physiology of the Eye*. 5th Edition. Pergamon Press: New York, NY.
- Dede, C., R.B. Loftin, and J.W. Regian. 1994. "The Design of Artificial Realities to Improve Learning Newtonian Mechanics." In *Proc. East-West International Conference on Multimedia, Hypermedia, and Virtual Reality*, Moscow, Russia, September 14-16.
- Durlach, N. I. and A.S. Mavor (Eds). 1995. *Virtual Reality Scientific and Technological Challenges*. National Academy Press: Washington, DC.
- Eggleston, R.G., W.P. Janson, and S. Adapalli. 1993. "Manual Tracking Performance Using a Virtual Hand Controller: A Comparison Study." In *AGARD Conference Proceedings 541, Virtual Interfaces: Research and Applications*, pp. 10-1 to 10-7.

- Emura, S. and S. Tachi. 1994. "Sensor Fusion Measurement of Human Head Motion." In *Proc. IEEE International Workshop on Robot and Human Communication*, pp. 124-129.
- Fausti, S.A., R.H. Frey, D.A. Erickson, B.Z. Rappaport, E.J. Cleary, and R.E. Brummett. 1982. "A System for Evaluating Auditory Function from 8000 to 20000 Hz." *Journal of the Acoustical Society of America*, 66, pp. 1713-1718.
- Finch, M., V.L. Chi, R.M. Taylor, M. Falvo, S. Washburn, and R. Superfine. 1995. "Surface Modification Tools in a Virtual Environment Interface to a Scanning Probe Microscope." In *Proc. ACM Symposium on Interactive 3D Graphics*, Monterey, CA, April 9-12, pp. 13-18.
- Foxlin, F. and N. Durlach. 1994. "An Inertial Head-Orientation Tracker with Automatic Drift Compensation for Use with HMDs." In *Virtual Reality Software and Technology, Proceedings of the VRST '94 Conference*, Singapore, August 23-26, pp. 159-173.
- Fukui, Y. and M. Shimojo. 1994. "Recognition of Virtual Shape Using Visual and Tactile Sense under Optical Illusion." In *Proc. 3rd IEEE International Workshop on Robot and Human Communication*, Nagoya, Japan, July 18-20, pp. 294-298.
- Gomez, D., G. Burdea, and N. Langrana. 1995. "Integration of the Rutgers Master II in a Virtual Reality Simulation." In *Proc. IEEE VRAIS '95 International Symposium*, Research Triangle Park, NC, March 11-15, pp. 198-202.
- Gottschalk, S. and J.F. Hughes. 1993. "Autocalibration for Virtual Environments Tracking Hardware." In *Computer Graphics Proceedings, Annual Conference Series*, pp. 65-71.
- Greene, W.B. and J.D. Heckman. 1994. *The Clinical Measurement of Joint Motion*. American Academy of Orthopedic Surgeons: Rosemont, IL.
- Hancock, D. 1993. "Prototyping the Hubble Fix." *Spectrum*, 30(10), pp. 34-39.
- Hannaford, B. 1989. "A Design Framework for Teleoperators with Kinesthetic Feedback." *IEEE Trans. on Robotics and Automation*, 5(4), 426-434.
- Hasser, C.J. and M.W. Daniels. 1996a. "Tactile Feedback with Adaptive Controller for a Force-Reflecting Haptic Display. Part 2: Improvements and Evaluation." In *Proc. 15th Southern Biomedical Engineering Conference*, Dayton, OH, March 29-31.

- Hasser, C.J. March 1996b. Personal communication.
- Hasser, C.J. July 1995. *Force-Reflecting Anthropomorphic Hand Masters*. Crew Systems Directorate, Biodynamics and Biocommunications Division, Armstrong Laboratory. AL/CF-TR-1995-0110.
- Hasser, C.J. and J.M. Wesenberger. 1993. "Preliminary Evaluation of a Shape-Memory Alloy Tactile Feedback Display." In *DSC-Vol. 49, Advances in Robotics, Mechatronics, and Haptic Interfaces*, pp. 73-80.
- Hayward, V. 1995. "Towards a Seven Axis Haptic Device." In *Proc. IROS '95, IEEE/RSI International Conference on Intelligent Robots and Systems*, Pittsburgh, PA, August 5-9.
- Hendrix, C.M. 1995. *Exploratory Studies on the Sense of Presence in Virtual Environments as a Function of Visual and Auditory Display Parameters*. Available on the WWW at URL <http://www.hitl.washington.edu>.
- Hill, J.W. and J.K. Salisbury, Jr. November 1976. *Study to Design and Develop Remote Manipulator Systems*. SRI Project 4055. SRI International, Menlo Park, CA.
- Hinchcliffe, R. 1959. "The Threshold of Hearing as a Function of Age." *Acoustica*, 9, pp. 303-308.
- Howe, R.D., W.J. Peine, D.A. Kontarinis, and J.S. Son. November 1994. "Remote Palpation Technology for Surgical Applications." Submitted to *IEEE Engineering in Medicine and Biology Magazine*.
- Howe, R.D. 1992. "A Force-Reflecting Teleoperated Hand System for the Study of Tactile Sensing in Precision Manipulation." In *Proc. IEEE Conference on Robotics and Automation*, Nice, France, May 12-14, pp. 1321-1326.
- Huang, J. and L.R. Young. 1981. "Sensation of Rotation About a Vertical Axis with a Fixed Visual Field in Different Illuminations and in the Dark." *Experimental Brain Research*, 41(2), pp. 172-183.
- Hubel, D. H. 1963. "The Visual Cortex of the Brain." *Scientific American*, 209(5).
- Hunter, I.W., T.D. Doukoglou, S.R. Lafontaine, P.G. Charette, L.A. Jones, M.A. Sagar, G.D. Mallinson, and P.J. Hunter. 1993. "A Teleoperated Microsurgical Robot and Associated Virtual Environment for Eye Surgery." *Presence*, 2(4), pp. 265-280.

- Ino, S., S. Shimizu, T. Odagawa, M. Sato, M. Takahashi, T. Izumi, and T. Ifukube. 1993. "A Tactile Display for Presenting Quality of Materials by Changing the Temperature of Skin Surfaces." In *Proc. IEEE 2nd International Workshop on Robot and Human Communication*, Tokyo, Japan, November 3-5, pp. 220-224.
- INTERACT. 1995. *The Gesture and Speech Integration Project*. Available on the WWW at URL <http://www.cs.cmu.edu/Web/Groups/interact/cal/cal.html>.
- International Standards Organization. 1975. *Standard Reference Zero for the Calibration of Pure-Tone Audiometers 389*. ISO: New York, NY.
- International Standards Organization. 1961. *Normal Equal-Loudness Contours for Pure Tones and Normal Threshold of Hearing under Free Field Listening Conditions R226*. ISO: New York, NY.
- Ishii, M. and M. Sato. 1994a. "A 3D Spatial Interface Device Using Tensed Strings." *Presence*, 3(1), pp. 81-86.
- Ishii, M. and M. Sato. 1994b. "Force Sensations in Pick-and-Place Tasks." In *DSC-Vol. 55-1, Dynamic Systems and Control: Volume 1*, pp. 339-344.
- Iwata, H. and K. Matsuda. 1992. "Haptic Walkthrough Simulator: Its Design and Application to Studies on Cognitive Map." In *Proc. International Conference on Artificial Reality and Tele-Existence*, pp. 185-192.
- Jacob, R.J.K. 1995. "Eye Tracking in Advanced Interface Design." In W. Barfield and T.A. Furness (Eds), *Virtual Environments and Advanced Interface Design*, pp. 258-288. Oxford University Press: New York, NY.
- Jacobus, H.N., A.J. Riggs, C.J. Jacobus, and Y. Weinstein. 1992. "Implementation Issues for Telerobotic Handcontrollers: Human-Robot Ergonomics." In *Human-Robot Interaction*, pp. 284-314. Taylor and Francis: Washington, DC.
- Jewell, M. R, G.R. Chamberlin, D.D. Sheat, P. Cochrane, and D.J. McCartney. 1995. "3-D Imaging Systems for Video Communication Applications." In *Proc. SPIE Vol. 2409, Stereoscopic Displays and Virtual Reality II*, pp 4-10.
- Johansson, R.S. 1991. "How is Grasping Modified by Somatosensory Input?" In D.R. Humphrey and H.-J Freund (Eds), *Motor Control: Concepts and Issues*, pp. 331-355. John Wiley and Sons: New York, NY.

- Kalawsky, R. S. 1993. *The Science of Virtual Reality and Virtual Environments: A Technical, Scientific and Engineering Reference on Virtual Environments*. Addison-Wesley Publishing Company: Reading, MA.
- Kessler, G.D., L.F. Hodges, and N. Walker. 1994. *Evaluation of the CyberGlove as a Whole Hand Input Device*. Georgia Institute of Technology.
- Knasko, S.C. and A.N. Gilbert. 1990. "Emotional State, Physical Well-Being, and Performance in the Presence of Feigned Ambient Odor." *Journal of Applied Social Psychology*, 20(16), pp. 1345-1357.
- Kocian, D.F. and H.L. Task. 1995. "Visually Coupled Systems Hardware and the Human Interface." In W. Barfield and T.A. Furness (Eds), *Virtual Environments and Advanced Interface Design*, pp. 175-257. Oxford University Press: New York, NY.
- Kontarinis, D.A. and R.D. Howe. 1995. "Tactile Display of Vibratory Information in Teleoperation and Virtual Environments." To appear in *Presence*.
- Kosslyn, S. E. 1994. *Image and Brain: The resolution of the imagery debate*. MIT Press: Cambridge, MA.
- Kosslyn, S. E. 1980. *Image and Mind*. Harvard University Press: Cambridge, MA.
- Krieg, J.C. 1992. "Accuracy, Resolution, Latency and Speed: Key Factors in Virtual Reality Tracking Environments." In *Proc. Virtual Reality '92, 3rd Annual Conference and Exhibition*, San Jose, CA, September, pp. 90-100.
- Krueger, M.W. 1996. Personal communication.
- Krueger, M.W. 1995. *Olfactory Stimuli in Virtual Reality Medical Training*. Artificial Reality Corporation, Vernon, CT.
- Laing, D.G. 1986. "Identification of Single Dissimilar Odors is Achieved by Humans with a Single Sniff." *Physiology and Behavior*, 37, pp. 163-170.
- LaMotte, R.H. 1978. "Intensive and Temporal Determinants of Thermal Pain." In D.R. Kenshalo (Ed.), *Sensory Functions of the Skin of Humans*. Plenum Press: NY.
- Landis, C. 1954. "Determinants of the Critical Flicker-Fusion Threshold." *Physiological Review*, 34, 259-286.
- Lipscomb, J. S. and W.L. Wooten. 1994. "Reducing Crosstalk Between Stereoscopic Views." In *SPIE Vol. 2177*, pp. 92-95.

- Magee, L.E. 1995. Personal communication.
- Mandelbaum, J. and L.L. Sloan. 1947. "Peripheral Visual Acuity." *American Journal of Ophthalmology*, 30, pp. 581-588.
- Massie, T.H. and J.K. Salisbury. 1994. "The PHANTOM Haptic Interface: A Device for Probing Virtual Objects." In *DSC, Vol. 55-1, Dynamic Systems and Control: Vol. 1*, pp. 295-299.
- Massimino, M. and T. Sheridan. 1993. "Sensory Substitution for Force Feedback in Teleoperation." *Presence* 2(4), pp. 344-352.
- McKenna, M. and D. Zeltzer. 1992. "Three Dimensional Visual Display Systems for Virtual Environments." *Presence*, 1(4) pp. 421-458.
- Meyers, K., H. Applewhite, and F.A. Biocca. 1992. "A Survey of Position Trackers." *Presence*, 1(2), pp. 173-200.
- Millman, P. and J.E. Colgate. 1994. "Experimental Study of Haptic Information Pickup in Manipulation." In *Proc. SPIE, International Symposium on Photonics for Industrial Applications, Telemanipulator and Telepresence Technologies Conference*, Boston MA.
- Moller, A.R. 1983. "Changes in hearing measures with increasing age." In R. Hinchcliffe (Ed.), *Hearing and Balance in the Elderly*. Churchill-Livingstone: Edinburgh, UK.
- Morgan, C.T. 1965. *Physiological Psychology*. McGraw-Hill, Inc.: New York, NY.
- Napier, J.R. 1956. "The Prehensile Movements of the Human Hand." *Journal of Bone and Joint Surgery*, 38B(4), pp. 902-913.
- Naus, A. 1985. "The Occupational Meaning of Smell." *Journal of Hygiene, Epidemiology, Microbiology, and Immunology*, 29(1), pp. 29-36.
- Ouh-Young, M., D.V. Beard, and F.P. Brooks, Jr. 1989. "Force Display Performs Better than Visual Display in a Simple 6-D Docking Task." In *Proc. IEEE International Conference on Robotics and Automation*, May 14-19, pp. 1462-1466.
- Overbosch, P., R. de Wijke, Th. J.R. de Jonge, and E.P. Koster. 1989. "Temporal Integration and Reaction Times in Human Smell." *Physiology and Behavior*, 45, pp. 615-626.

- Patrick, N.M., T.B. Sheridan, and M.J. Massimino. 1990. "Design and Testing of a Non-Reactive, Fingertip, Tactile Display for Interaction with Remote Environments." In *Proc. SPIE Vol. 1837, Cooperative Intelligent Robotics in Space*, pp. 215-222.
- Pausch, P., M.A. Shackelford, and D. Proffitt. 1993. "A User Study Comparing Head-Mounted and Stationary Displays." In *Proc. IEEE Symposium on Research Frontiers in Virtual Reality*, pp. 41-45.
- Peters, R.A. April 1969. *Dynamics of the Vestibular System and Their Relation to Motion Perception, Spatial Disorientation, and Illusions*. NASA CR-1309, NASA.
- Reynier, F. and V. Hayward. March 1993. *Summary of the Kinesthetic and Tactile Function of the Human Upper Extremities*. McGill Research Center for Intelligent Machines, McGill University, Montreal, Canada.
- Richards, A. November 1994. *Stroboscopic Display Unit*. International Patent Application published by the World Intellectual Property Organization, No. WO 94/25899.
- Robinson, D.W., M.S. Shipton, and R. Hinchcliffe. 1981. "Audiometric Zero for Air Conduction." *Audiology*, 20, pp. 409-431.
- Robinson, D.W. and G.J. Sutton. 1979. "Age Effect in Hearing—A Comparative Analysis of Published Threshold Data." *Audiology*, 18, pp. 320-324.
- Robinson, D.W. and R.S. Dadson. 1956. "A Re-Determination of Equal-Loudness Relations for Pure Tones." *British Journal of Applied Physics*, 7, pp. 166-181.
- Scharf, B. and S. Buus. 1986. "Audition I: Stimulus, Physiology, Thresholds." In K.R. Boff, L. Kaufman, and J.P. Thomas (Eds), *Handbook of Perception and Human Performance, Volume I: Sensory Process and Perception*. John Wiley and Sons: New York, NY.
- Schlesinger, G. 1919. Der Mechanische Aufbau der Kuntstlichen Glieder. In M. Borchardt et al (Eds), *Ersatzglieder und Arbeitshilfen fur Kriegsbeschadigte und Unfallverletzte*. pp. 321-699. Springer-Verlag: Berlin, Germany.
- Schneider, W. 1988. "The Tactile Array Stimulator." *John Hopkins APL Technical Digest*, 9(1), pp. 39-43.
- Segal, N.L., T.D. Topolski, S.M. Wilson, K.W. Brown, and L. Araki. 1995. "Twin Analysis of Odor Identification and Perception." *Physiology and Behavior*, 57(3), pp. 605-609.

- Sherrick, C.E. 1984. "Basic and Applied Research on Tactile Aids for Deaf People: Progress and Prospects." *Journal of the Acoustical Society of America*, 75(5), pp. 1325-1342.
- Shimizu, Y. 1986. "Tactile Display Terminal for Visually Handicapped." *Displays*, July, pp. 116-120.
- Shimoga, K.B. 1993a. "A Survey of Perceptual Feedback Issues in Dextrous Telemanipulation: Part I. Finger Force Feedback." In *Proc. IEEE Annual Virtual Reality International Symposium*, Seattle, Washington, September 18-22, pp. 263-270.
- Shimoga, K.B. 1993b. "A Survey of Perceptual Feedback Issues in Dextrous Telemanipulation: Part II. Finger Touch Feedback." In *Proc. IEEE Annual Virtual Reality International Symposium*, Seattle, Washington, September 18-22, pp. 271-279.
- Slater, M., M. Usoh, and A. Steed. 1994. "Steps and Ladders in Virtual Reality." In *Proc. VRST '94, Virtual Reality Software and Technology*, Singapore, 23-26 August, pp. 45-54.
- Stelmachowicz, P.G., M.P. Gorga, and J.K. Cullen. 1982. "A Calibration Procedure for the Assessment of Thresholds Above 8000 Hz." *Journal of Speech and Hearing Research*, 25, pp. 681-683.
- Strickland, D., A. Patel, C. Stovall, J. Palmer, and D. McAllister. 1994. "Self Tracking of Human Motion for Virtual Reality Systems. In *SPIE Vol. 2177*, pp. 278-287.
- Sturman, D.J. and D. Zeltzer. 1994. "A Survey of Glove-Based Input." *IEEE Computer Graphics and Applications*, January, pp. 30-39.
- Sturman, D.J. and D. Zeltzer. 1993. "Utility of Whole-Hand Input." In *Proc. SPIE Vol. 2057*, pp. 282-291.
- Summer, S. K. and B. Katz. 1994. *Device for the Creation of Three-Dimensional Images*. U. S. Patent No. 5,311,357.
- Sutter, P.H., J.C. Iatridis, and N.V. Thakor. 1989. "Response to Reflected-Force Feedback to Fingers in Teleoperations." In *Proc. NASA Conference on Space Telerobotics*, NASA JPL.
- Tan, H.Z., M.A. Srinivasan, B. Eberman, and B. Cheng. 1994. "Human Factors for the Design of Force-Reflecting Haptic Interfaces. In *DSC, Vol. 55-1, Dynamic Systems and Control: Vol. 1*. pp. 353-359.

- Tanner, S. 1993. "The Use of Virtual Reality at Boeing's Huntsville Laboratories." In *Proc. 1993 IEEE Virtual Reality Annual International Symposium*, September 18-22, pp. 14-19.
- Tecumeh, W. and G. Kramer. 1994. "Sonifying the Body Electric: Superiority of an Auditory over a Visual Display in a Complex, Multivariate System." In *Santa Fe Institute Studies in the Science of Complexity, Proc. Vol. XVIII*, pp. 307-325. Addison-Wesley: Reading, MA.
- Wang, J.-F., R. Azuma, G. Bishop, V. Chi, J. Eyles, and H. Fuchs. 1990. "Tracking a Head-Mounted Display in a Room-Sized Environment with Head-Mounted Camera." In *SPIE, Technical Symposium on Optical Engineering and Photooptics in Aerospace Sensing*, Orlando, Florida, April.
- Ware, C., C. Gobrecht, and M. Paton. 1995. "Algorithm for Dynamic Disparity Adjustment." In *Proc. SPIE Vol. 2409, Stereoscopic Displays and Virtual Reality II*, pp. 150-156.
- Ware, C. and R. Balakrishnan. December 1994. "Reaching for Objects in VR Displays: Lag and Frame Rate." *ACM Transactions on Computer-Human Interaction*, 1(4), pp. 331-356.
- Ware, C., K. Arthus, and K.S. Booth. 1993. "Fish Tank Virtual Reality. In *Interchi '93: Bridges between Worlds*, pp. 37-42. Addison-Wesley: Reading, MA.
- Weissler, P.B. 1968. "International Standard Reference Zero for Audiometers." *Journal of the Acoustical Society of America*, 44, pp. 264-275.
- Wenzel, E.M. 1994. "Spatial Sound and Sonification." In G. Kramer (Ed.), *Santa Fe Institute Studies in the Sciences of Complexity, Proceedings Volume XVIII*. Addison-Wesley: New York, NY.
- Wenzel, E.M., M. Arruca, D.J. Kistler, and F.L. Wightman. 1993. "Localization Using Nonindividualized Head-Related Transfer Functions." *Journal of the Acoustical Society of America*, 94(1), pp. 111-123.
- Wenzel, E.M. 1992. "Localization in Virtual Acoustic Displays." *Presence*, 1(1), pp. 80-107.
- Wiker, S.F., E. Hershkowitz, and J. Zilk. 1989. "Teleoperator Comfort and Psychometric Stability: Criteria for Limiting Master-Controller Forces of Operation and Feed-

back During Telemanipulation." In *Proc. NASA Conference on Space Telerobotics*, Vol. I, Pasadena, CA, January/February, pp. 99-107.

Wise, S., J. Rosen, W. Gardner, E. Sabelman, E. Valainis, Y. Wong, K. Glass, J. Drace, and J. Rosen. 1990. "Evaluation of a Fiber-Optic Glove for Semi-Automated Goniometric Measurements." *Journal of Rehabilitation Research and Development*, 27(4), pp. 411-424.

Yarnitsky, D. and J.L. Ochoa. 1991. "Warm and Cold Specific Somatosensory Systems." *Brain*, 114, pp. 1819-1826.

Zellner, D.A., A.M. Bartoli, and R. Eckard. 1991. "Influence of Color on Odor Identification and Liking Ratings." *American Journal of Psychology*, 104(4), pp. 547-561.

Zerkus, M., B. Becker, J. Ward, L. Halvorsen. *Thermal Feedback in Virtual Reality and Telerobotic Systems*. Technical Report N95-15987. NASA Scientific and Technical Information Office, NASA.

Zhai, S. and P. Milgram. 1993. "Human Performance Evaluation of Manipulation Schemes in Virtual Environments." In *Proc. 1st Virtual Reality Annual International Symposium (VRAIS)*, Seattle, pp. 155-161.

Zhai, S. 1993. "Investigation of Feel for 6DOF Inputs: Isometric and Elastic Rate Control for Manipulation in 3D Environments." In *Proc. Human Factors and Ergonomics Society 37th Annual Meeting*, pp. 323-327.

LIST OF ACRONYMS AND ABBREVIATIONS

2-D	Two-Dimensional
3-D	Three-Dimensional
AIST	Agency of Industrial Science and Technology
ALIVE	Artificial Life Interaction Video Environment
ARC	Artificial Reality Corporation
ARL	Army Research Laboratory
ARM	Argonne Remote Manipulator
ARPA	Advanced Research Projects Agency
CAD	Computer-Aided Design
CCD	Charge Coupled Device
CGSD	Computer Graphics Systems Development
CRT	Cathode Ray Tube
DARPA	Defense Advanced Research Projects Agency
DIBBL	Dismounted Infantry Battlespace BattleLab
DDL	Denne Development Limited
DIS	Distributed Interactive Simulation
DNA	Deoxyribonucleic acid
DoD	Department of Defense
DPI	Dual-Purkinje-Image
DSP	Digital Signal Processing
DTSS	Displaced Temperature Sensing System
DOF	Degree of Freedom
EAM	Exoskeleton ArmMaster
EOG	Electrooculogram
EPI	Epipolar-Plane Image
EVA	Extra Vehicular Activity
FTTE	Fusion Interfaces for Tactical Environments
FTFS	Force and Tactile Feedback System
GUI	Graphical User Interface

HAPTAC	HAPtic-TACtile
HDVD	High Definition Volumetric Display
HEHD	Hand Exoskeleton Haptic Display
HHT	Head Hand XYZ Tracker
HITL	Human Interface Technology Laboratory
HMD	Head-Mounted Display
HRTF	Head-Related Transfer Function
IID	Interaural Intensity Difference
IIT	Interaural Time Difference
INTERACT	Interactive Systems Laboratory
IPD	Inter-Pupil Distance
IPORT	Individual Portal
IR	Infra-Red
ISMS	Intelligent Soldier Mobility Simulator
JND	Just-Noticeable-Difference
JPL	Jet Propulsion Laboratory
LCD	Liquid Crystal Diode
LED	Light Emitting Diode
MAF	Minimum Audible Field
MAP	Minimum Audible Pressure
MCP	Metacarpophalangeal
MIT	Massachusetts Institute of Technology
MITI	Ministry of International Trade and Industry
MSR	Microsurgical Robot
nM	Nanomanipulator
NASA	National Aeronautics and Space Administration
NAWCTSD	Navy Air Warfare Center Training Systems Division
NPSNET	Naval Postgraduate School Networked Vehicle Simulator
NSF	National Science Foundation
PAS	Polyphonic Audio Spatializer
PDMFF	Portable Dextrous Master with Force Feedback
PemRAM	Precision Electromagnetic RAM
PIP	Proximal Interphalangeal
PMP	Personal Motion Platform
PUSH	Personal Use Stereoscopic Haptic

RM	Rutgers Master
SBIR	Small Business Innovation Research
SMA	Shape Memory Alloy
SMART	Simultaneous Multiple Area Recognition and Tracking
SPIDAR	Spatial Interface Device for Artificial Reality
STRICOM	Simulation Training and Instrumentation Command
TRP	Technology Reinvestment Program
UNC	University of North Carolina
US	United States
VAPS	Virtual Audio Processing System
VE	Virtual Environment
VEIL	Virtual Environment Interface Laboratory
VR-B	Virtual Binoculars
WPAFB	Wright Patterson Air Force Base

APPENDIX A. POINTS OF CONTACT

A.1 Vendors

A.1.1 Visual Interfaces

Bogdanski Advanced Video Systems AB
Hogsommarvagen 11
S-142 42 Skogas, Sweden
Telephone: +46 8 448 20 25
Fax: +46 8 448 20 26

Dimension Technologies, Inc.
315 Mount Read Boulevard
Rochester, NY 14611
Telephone: 716-436-3530
Fax: 716-436-3280

Ericsson Microway Systems AB
Airborne Electronics Division
S-16484 Stockholm, Sweden
Telephone: +46 8 757 30 00
Fax: +46 8 752 81 72

Fakespace, Inc.
4085 Campbell Avenue
Menlo Park, CA 94025
Telephone: 415-688-1940
Fax: 415-688-1949

Forte Technologies, Inc.
2615 W. Heneretta Road
Rochester, New York 14623
Telephone: 716-427-8604
Fax: 716-427-7383

General Reality Company
124 Race Street
San Jose, CA 95126
Telephone: 408-289-8340
Fax: 408-289-8258

Kaiser Electro-Optics, Inc.
2752 Loker Avenue
Carlsbad, CA 92008
Telephone: 619-438-9255
Fax: 619-438-6875

Leep Systems, Inc.
241 Crescent Street
Waltham, MA 02154
Telephone: 617-647-1395
Fax: 617-647-1109

Liquid Image Corporation
659 Century Street
Winnipeg, Manitoba, R3H0L9, Canada
Telephone: 204-775-2633
Fax: 204-772-0239

MicroSharp

N-Vision, Inc.
7680 Old Springhouse Road
Madison Building, 1st Floor
McLean, VA 22102
Telephone: 703-506-8808
Fax: 703-903-0455

Phoenix Group, Inc.
204 Terminal Drive
Plainview, NY 11803
Telephone: 516-349-1919
Fax: 516-349-1926

Victor Maxx Technologies
510 Lake Cook Road, Suite 100
Deerfield, IL 60015
Telephone: 708-267-0007
Fax: 708-267-0034

Virtual I/O, Inc.
1000 Lenora Street, Suite 600
Seattle, WA 98121
Telephone: 206-382-7410
Fax: 206-382-8810

Virtual Reality, Inc.
800 Follin Lane, Suite 270
Vienna, VA 22180
Telephone: 703-242-0030
Fax: 703-242-5220

Virtual Research Systems, Inc.
3193 Belick Street, #2
Santa Clara, CA 95054
Telephone: 408-748-8712
Fax: 408-748-8714

VRex, Inc.
8 Skyline Drive
Hawthorne, NY 10532
Telephone: 914-345-8877
Fax: 914-345-9558

A.1.2 Tracking Interfaces

Ascension Technology Corporation
P.O. Box 527
Burlington, VT 05402
Telephone: 802-860-6440
Fax: 802-860-6439
E-mail: ascension@world.std.com

BioControl Systems, Inc.
2555 Park Boulevard
Palo Alto, CA 94306
Telephone: 415-329-8494
Fax: 415-329-8498
E-mail: biomuse@well.com

BioVision
Optimum Human Performance Center
1580 California Street
San Francisco, CA 94109
Telephone: 415-292-0333
Fax: 415-292-0344
E-mail: <http://www.bio-vision.com>

CCG MetaMedia, Inc.
460 West 34th Street
New York 10001
Telephone: 212-268-2100
Fax: 212-268-2105

FakeSpace, Inc.
4085 Campbell Avenue
Menlo Park, CA 94025
Telephone: 415-688-1940
Fax: 415-688-1949

Fifth Dimension Technologies
P.O. Box 5, Persequor Technopark
Persequor 0020
Pretoria, South Africa
Telephone: +27-12349-1400
Fax: +27-12349-1404
E-mail: 5dt@pixie.co.za

Forward Technologies
777 Alvarado Road, Suite 266
La Mesa, CA 91941
Telephone: 619-258-8789
Fax: 619-258-8840

General Reality Company
124 Race Street
San Jose, CA 95126
Telephone: 408-289-8340
Fax: 408-289-8258
E-Mail: GRCsales@aol.com

Gyration, Inc.
12930 Saratoga Avenue, Building C
Saratoga, CA 95070
Telephone: 408-255-3016
Fax: 408-255-9075

Hughes Training Corporation
Link Division
P.O. Box 1237
Binghamton, NY 13902-1237
Telephone: 607-721-4604
Fax: 607-721-5600

Innovision Systems
30521 Schoenherr, Suite 104
Warren, MI 48093-3129
Telephone: 810-751-0646
Fax: 810-751-0600

ISCAN, Inc.
P.O. Box 382076
Cambridge, MA 02238-2076
Telephone: 617-273-4455
Fax: 617-273-0076
E-mail: iscan@cambridge.village.com

LC Technologies, Inc.
9455 Silver King Court
Fairfax, VA 22031
Telephone: 703-385-7133
Fax: 703-385-7137
E-mail: <http://www.lctinc.com>

Logitech, Inc.
6607 Kaiser Drive
Fremont, CA 94555
Telephone: 510-713-4783
Fax: 510-505-0979

Northern Digital, Inc.
403 Albert Street
Waterloo
Ontario, N2L 3V2, Canada
Telephone: 519-884-5142
Fax: 519-884-5184

Origin Instruments
854 Greenview Drive
Grand Prairie, TX 75050-2438
Telephone: 214-606-8740
Fax: 214-606-8741

Polhemus
P.O. Box 560
Colchester, VT 05446
Telephone: 802-655-3159
Fax: 802-655-1439

Precision Navigation
1235 Pear Avenue, Suite 111
Mountain View, CA 94043
Telephone: 415-962-8777
Fax: 415-962-8776

RPI Advanced Technology Group
P.O. Box 14607
San Francisco, CA 94114
Telephone: 415-495-5671
Fax: 415-495-5124

Science Accessories
7125 Riverwood Drive
Columbia, MD 21046
Telephone: 410-381-6688
Fax: 410-290-9065

Shooting Star Technology
52023 Yale Road, RR1,
Rosedale B.C., Canada VOX 1X
Telephone: 604-794-3364
Fax: 604-794-3139

SMX Corporation
222 Gale Lane
Kennett Square, PA 19348
Telephone: 610-444-2300
Fax: 610-444-2323

Systron Donner Inertial Division
2700 Systron Drive
Concord, CA 94518-1399
Telephone: 510-671-6648
Fax: 510-671-6647
E-mail: service@systron.com

The Vivid Group
317 Adelaide Street West, Suite 302
Toronto, Ontario, M5V-1P9, Canada
Telephone: 416-340-9290
Fax: 416-348-9808

VictorMaxx Technologies
510 Lake Cook Road, Suite 100
Deerfield, IL 60015
Telephone: 708-267-0007
Fax: 708-267-0037

Vidtronics Inc.
4610 Delemere Boulevard
Royal Oak, MI 48073
Telephone: 810-549-7444
Fax: 810-549-7687

A.1.3 Auditory Interfaces

Audio Cybernetics
655 Highway 72
Golden, CO 80403
Telephone/Fax: 303-642-0164
E-mail: audiocy@rmii.com

Crystal River Engineering
Professional Audio Group
490 California Avenue, Suite 200
Palo Alto, CA 94306
Telephone: 415-323-8155
Fax: 415-963-9188
E-mail: proaudio@cre.com

Focal Point Three-Dimensional Audio
1402 Pine Avenue, Suite 127
Niagara Falls, NY 14301
Telephone: 416-963-9188
Fax: 416-963-9188
E-mail: gehring@dgp.toronto.edu

QSound Labs
2748-37 Avenue, NE
Telephone: 403-291-2492
Fax: 403-250-1521
E-mail: 75033,53@compuserve.com

Reality by Design

Roland Corporation US
7200 Dominion Circle
Los Angeles, California 90040-3696
Telephone: 213-685-5141
Fax: 213-726-8865

A.1.4 Primary User Input Interfaces

EXOS, Inc.
2A Gill Street
Woburn, MA 01801
Telephone: 617-933-0022
Fax: 617-933-0303

Fakespace, Inc.
4085 Campbell Avenue
Menlo Park, CA 94025
Telephone: 415-688-1940
Fax: 415-688-1949

Fifth Dimension Technologies
P.O. Box 5, Perseus Technopark
Perseus 0020
Pretoria, South Africa
Telephone: +27-12349-1400
Fax: +27-12349-1404

InWorld VR, Inc.
3030 Bridgeway #207
Sausalito, CA 94965
Telephone: 415-331-5004
Fax: 415-331-5010
E-mail: inworld@well.com

Immersion Corporation
2158 Paragon Drive
San Jose, CA 95131
Telephone: 408-467-1900
E-mail: immersion@starconn.com

Kantek, Inc.
1414 Avenue of the Americas
New York, NY 10019
Telephone: 212-355-5049
Fax: 212-593-0065

Logitech, Inc.
6607 Kaiser Drive
Fremont, CA 94555
Telephone: 800-889-0065
Fax: 510-505-0979

Spacetec IMC Corporation
600 Suffolk Street
Lowell, MA 01854-3629
Telephone: 508-970-0330
Fax: 508-970-0199

T.C.A.S. Effects, Ltd.
130 City Road
Cardiff CF2 3DR, UK
Telephone: +44-0-1222-489888
Fax: +44-0-1222-619770

Virtual Technologies, Inc.
2175 Park Blvd.
Palo Alto, CA 94306
Telephone: 415-321-4900
Fax: 415-321-4912
E-mail: info@virtex.com

A.1.5 Haptic Interfaces

Aura Systems, Inc.
2335 Alaska Avenue
El Segundo, CA 90245
Telephone: 800-909-aura
Fax: 310-643-9463

CM Research, Inc.
2437 Bay Area Blvd.
Houston, TX 77058
Telephone: 713-488-3598
Fax: 713-488-3599

Cybernet Systems Corporation
1919 Green Road, Suite B-101
Ann Arbor, MI 48105
Telephone: 313-668-2567
Fax: 313-668-8780

EXOS, Inc.
2A Gill Street
Woburn, MA 01801
Telephone: 617-933-0022
Fax: 617-933-0303

Immersion Corporation
2158 Paragon Drive
San Jose, CA 95131
Telephone: 408-467-1900
E-mail: immersion@starconn.com

Nissho Electronic (USA) Corporation
18201 Von Karman Avenue, Suite 350
Irvine, CA 92715
Telephone: 714-261-8811
Fax: 714-261-8819

SensAble Technologies
26 Landsdowne Street
University Park at MIT
Cambridge, MA 02139
Telephone: 617-621-0150
Fax: 617-621-0135
E-mail: sensable@sensable.com

Virtual Technologies, Inc.
2175 Park Blvd.
Palo Alto, CA 94306
Telephone: 415-321-4900
Fax: 415-321-4912

Xtensory, Inc.
140 Sunridge Drive
Scotts Valley, CA 95066
Telephone: 408-439-0600
Fax: 408-439-8845
E-mail: cutt@netcom.com

A.1.6 Full Body Motion Interfaces

Aerotrim USA, Inc.
17437 SE Kendall Court
Portland, OR 97236
Telephone: 503-665-7800
Fax: 503-665-0472

CineMotion International plc
Oak Court
7 Betts Way, London Road
Crawley, West Sussex RH10 2GB, UK
Telephone: +44-0-1293-519192
Fax: +44-0-1293-519193

Denne Developments, Ltd.
Unit 4, Cedar Park
Cobham Road
Ferndown Industrial Estate
Wimborne, Dorset BH21 7SB, UK
Telephone: +44-1-202-861661
Fax: +44-2-793-432625

Dreamality Technologies, Inc.

Jessler Enterprises, Inc.
849 Fife's Bay Marina Road
Peterborough
Ontario, Canada K9J 6XS
Telephone: 705-876-0226
Fax: 705-876-1363

Orbotron, Inc.
2588 El Cumino Real
Carlsbad, CA 92008
POC: Chris Altaire
Telephone: 619-757-6900
Fax: 619-757-6996
E-mail: altare@ibm.net

RPI Entertainment
P.O. Box 14607
San Francisco, CA 94114
Telephone: 415-495-4460
Fax: 415-495-5124

StrayLight Corporation
150 Mount Bethel Road
Warren, NJ 07059
Telephone: 908-580-0086
Fax: 908-580-0092

Torus Systems, Inc.
3 Park Hill, Suite 4
Bickley, Kent BR1 2JH, UK
Telephone: +44-0-181-467-3651

Trailcraft Manufacturing, Ltd.
100 Wilkinson Road, Units 16-19
Brampton
Ontario, L6T 4Y9, Canada
Telephone: 905-450-1144
Fax: 905-450-6327

ViRtogo, Inc.
2222 North Elston Avenue
Chicago, IL 60614
Telephone: 312-281-6435
Fax: 312-281-1567

Virtual Images, Inc.
4388 Tuller Road
Dublin, OH 43017
Telephone: 614-799-1100
Fax: 614-799-1177

A.1.7 Olfactory Interfaces

BOC Group plc
The Priestly Centre
The Surrey Research Park
Guilford, Surrey GU2 5XY, UK
Telephone: +44-0-1483-450392
Fax: +44-0-1483-579857

Ferris Productions, Inc.
4001 E. Broadway, Suite 13
Phoenix, AR 85040
Telephone: 602-470-1177
Fax: 602-470-1166

A.2 Researchers

A.2.1 Visual Interfaces

British Aerospace plc
Sowerly Research Clinic, FPC 267
P.O. Box 5
Filton, Bristol BS12 7QW England
POC: Graham Edgar
Telephone: +44-117-969-3831
Fax: +44-117-036-3733

BT Laboratories
Martlesham Heath
Ipswich, UK
POC: Michael Jewell or Dennis Sheot
Telephone: +44-1473-643210
Fax: +44-1473-646885

Canon, Inc.
Media Technology Laboratory
Kawaski
Kanagawa 211, Japan
POC: Akihiro Katayama
Telephone: +81-3-334-82121

Dimensional Media Associates
675 Avenue of the Americas
New York, New York 10010
POC: Susan Kasen Summer
Telephone: 212-620-4100
Fax: 212-620-7771
E-mail: DMA@3DMedia.com

Dimension Technologies, Inc.
315 Mount Read Boulevard
Rochester, NY 14611
POC: Arnold D. Lagergren
Telephone: 716-426-3530
Fax: 716-436-3280

Electronic Visualization Laboratory
University of Illinois
Chicago, Illinois
POC: Carolina Cruz-Neira

Georgia Institute of Technology
Graphics Visualization and Usability Center
Atlanta, GA
POC: Wayne L. Wooten
Telephone: 404-894-3172
E-mail: wlw@cc.gatech.edu

IBM Thomas J. Watson Research Center
P.O. Box 704
Yorktown Heights, NY 10598
POC: James S. Lipscomb
Telephone: 914-784-7048
E-mail: lips@watson.ibm.com

Infinity Multimedia
14225 Ventura Boulevard
Sherman Oaks, CA 91423
POC: Peter M. Canepa
Fax: 818-501-6224

Perdue University
Department of Engineering Center
West Lafayette, Indiana 47907-1285
POC: Dr. Zygmunt Pizlo
Telephone: 317-494-6930
Fax: 317-496-1264

University of New Brunswick
Faculty of Computer Science
P.O. Box 4400
Fredericton, New Brunswick
Canada E3B 5A3
POC: Colin Ware
E-mail: cware@unb.ca

A.2.2 Tracking Interface

Artificial Reality, Inc.
P.O. Box 786
Vernon, CT 06066
POC: Dr. Myron Krueger
Telephone: 860-871-1375
Fax: 860-871-9445

Boeing Computer Services
Research and Technology
P.O. Box 24346, MS 7L-22
Seattle, WA 98124-0346
POC: David Mizzell and Adam Janin
Telephone: 206-865-3545
Fax: 206-865-2966
E-mail: janin@espresso.rt.cs.boeing.com

University of Washington
Human Interface Technology Laboratory
Seattle, WA 98195
POC: Michael Tidwell
Telephone: 206-616-1468
Fax: 206-543-5380
E-mail: tidwell@hitl.washington.edu

Xenotech
Suite 1, 41 Walters Drive
Osborne Park
Western Australia 6017
POC: Robert Baker
Telephone: +61 9 446 3366
Fax: +61-9-446-3340

Hughes Training-Link Corporation
Binghamton, NY
POC: David Peters
Telephone: 607-721-4604
Fax: 607-721-5600

Carnegie Mellon University
Interactive Systems Laboratories
School of Computer Science
5000 Forbes Avenue
Pittsburgh, PA 15213
POC: Alex Waibel
Telephone: 412-268-7676
Fax: 412-268-5578
E-mail: waibel@cs.cmu.edu

Computer Graphics Systems Development
Corporation
2483 Old Middlefield Way #140
Mountain View, CA 94043-2330
POC: Dr. Roy Latham
Telephone: 415-903-4922
Fax: 415-967-5252
E-mail: rlatham@cgsd.com

MIT Media Lab, Rm E15-305
20 Ames Street
Cambridge, MA 02139
POC: Pattie Maes
E-mail: pattie@media.mit.edu
<http://www.media.mit.edu>

NASA Ames Research Center
View Laboratory
MS T045-1
Moffett Field, CA 94035-1000
POC: Steve Bryson
Telephone: 415-604-4524
E-mail: bryson@nas.nasa.gov

Siemens AG
Corporate R&D Department ZFE ST SN 5
Otto Hahn Ring 6
81730 Munich, Germany
POC: Christopher Maggioni
Telephone: 49-89-636-42573
Fax: 49-89-636-2393
E-mail: Christoph.Maggioni@zfe.siemens.de

Sony Computer Science Laboratory, Inc.
Takanawa Muse Building
3-14-13 Higashi-gotanda
Shinagawa-ku, Tokyo 141, Japan
POC: Dr. Jun Rekimoto
Telephone: 81-3-5448-4380
Fax: 81-3-5448-4273
E-mail: rekimoto@csl.sony.co.jp

MIT
Research Laboratory of Electronics
Cambridge, MA 02139
POC: Nathaniel Durlach
Telephone: 617-253-2534
E-mail: durlach@cbgrrle.mit.edu

SUNY Stony Brook
Computer Science Department
Stony Brook, New York 11794-4400
POC: Professor Arie Kaufman
Telephone: 516-632-8470
Fax: 516-632-8334

University of North Carolina at Chapel Hill
Department of Computer Science CB#3175
Sitterson Hall
Chapel Hill, NC 27599
POC: Gary Bishop
Telephone: 919-962-1758
Fax: 919-962-1799
E-mail: houseman@cs.unc.edu

University of Washington
Human Interface Technology Laboratory
215 Fluke Hall
Box 352142
Seattle, WA 98195-2142
POC: Rich Johnston
Telephone: 206-685-3215
Fax: 206-543-5380
E-mail: richj@hitl.washington.edu

A.2.3 Auditory Interfaces

Graphics and Video Laboratory
Department of Computer Science
Naval Postgraduate School
Code CS/Zk
Monterey, California 93943
POC: Professor Michael J. Zyda
Telephone: 408-646-2305
Fax: 408-646-2814

Human Interface Research Branch
Aerospace Human Factors Research Division
NASA Ames Research Center
MS 262-2, Moffett Field
California 94035-1000
POC: Dr. Elizabeth M. Wenzel
Telephone: 415-604-6290
Fax: 415-604-3729

A.2.4 Primary User Input Interfaces

Armstrong Laboratory
Crew Systems Directorate
Virtual Environment Interface Laboratory
POC: Dr. Robert Eggleston
Telephone: 513-225-8764

Digital Image Design, Inc.
72 Spring Street
New York, NY 10012
POC: Juey Ong
Telephone: 212-343-2442
Fax: 212-343-0440

Department of Industrial Engineering
University of Toronto
POC: Dr. Shumin Zhai and Dr. Paul Milgram
Telephone: 416-978-3776
Fax: 416-978-3453
E-mail: zhai@ie.toronto.edu

Visualization and Usability Center
Georgia Institute of Technology
POC: Drew Kessler
Telephone: 404-853-9393
E-mail: drew@cc.gatech.edu

A.2.5 Haptic Interfaces

Begej Corporation
5 Claret Ash Road
Littleton, OH 80127-3536
POC: Steve Begej
Telephone: 303-932-2186
Fax: 303-932-2186
E-mail: li@ctsd2.jsc.nasa.gov

Computer Graphics Systems Development
Corporation
2483 Old Middlefield Way #140
Mountain View, CA 94043-2330
POC: Dr. Roy Latham
Telephone: 415-903-4922
Fax: 415-967-5252
E-mail: rlatham@cgsd.com

Boeing Computer Services
P.O. Box 24346, MS 7L-43
Seattle, WA 98124-0346
POC: Dr. William McNeely
Telephone: 206-865-3614
Fax: 206-865-2965
E-mail: wmcneely@espresso.rt.cs.boeing.com

Division of Applied Sciences
Harvard University
Pierce Hall
Cambridge, MA 02138
POC: Professor Robert Howe
Telephone: 617-496-8359
Fax: 617-495-9837

Hokkaido University
Research Institute for Electronic Science
Lab. of Sensory Information Engineering
Sapporo. 060, Japan
POC: Dr. Shuichi Ino
Telephone: +81-11-706-2414
Fax: +81-11-7-6-4968
E-mail: ino@sense.hokudai.ac.jp

McGill University
Research Center for Intelligent Machines
McConnell Engineering Building
3480 University Street
Montreal, Quebec, H3A 2A7, Canada
POC: Dr. Vincent Hayward
Telephone: 514-398-5006
Fax: 514-398-7348

MIT
Artificial Intelligence Laboratory
POC: Dr. Ken Salisbury
Telephone: 617-253-5834
Fax: 616-253-5060
E-mail: jks@ai.mit.edu

MIT
Department of Mechanical Engineering
POC: Dr. Ian Hunter
Telephone: 617-253-3921
E-mail: ihunter@mit.edu

MIT
Touch Laboratory
Department of Mechanical Engineering
POC: Dr. Mandayam Srinivasan
Telephone: 617-253-2512
Fax: 617-258-7003
E-mail: SRINI@cbgrle.mit.edu

Karlsruhe Research Center
Department of Engineering Technology
Abt. HIT
P.O. Box 3640
76021 Karlsruhe, Germany
POC: Dr. Harald Fischer
Telephone: +7247-823072
Fax: +87247-822289
E-mail: hafi@hdi.kfk.de

Rutgers University
Center for Computer Aids for Industrial Productivity
POC: Dr. Greg Burdea
Telephone: 908-445-5309
E-mail: burdea@telerobo.rutgers.edu

Sandia National Laboratories
P.O. Box 969, MS 9103
Livermore, CA 94551-0969
POC: Dave Andaleon
Telephone: 510-294-1552
Fax: 510-294-1377
E-mail: ddandal@sandia.gov

Suzuki Motor Corporation
Technical Research Center
2-1 Sakura-Namiki
Tsuzuki-Ku
Yokohama, 224, Japan
POC: Dr. Yoshitaka Adachi
Telephone: +81-45-943-7111
Fax: +81-45-943-7100

TiNi Alloy Company
1144 65th Street, Unit A
Oakland, CA 94608
Telephone: 510-483-9676
Fax: 510-483-1309

MITI/AIST
National Institute of Bioscience and Human
Technology
Japan
POC: Dr. Yukio Fukui
E-mail: fukui@nibh.go.jp

Northwestern University
Department of Mechanical Engineering
POC: Dr. Ed Colgate
E-mail: colgate@nwu.edu

University of North Carolina at Chapel Hill
Computer Science Department
POC: Linda Houseman
Telephone: 919-962-1758

University of Salford
Department of Electronic Engineering
POC: Dr. Darwin Caldwell
Telephone: +44-161-745-5000 ext. 4010
E-mail: b.d.caldwell@eee.salford.ac.uk

Tokyo Institute of Technology
Precision and Intelligent Laboratory
Japan
POC: Dr. Makoto Sato
Telephone: +81-45-924-5050
Fax: +81-45-921-0898

University of Hull
Department of Electrical Engineering
POC: Dr. Taylor
Telephone: +44-1-482-346311
Fax: +44-1-482-466664

University of Tsukuba
Institute of Engineering Mechanics
POC: Dr. Hiro Iwata
Telephone: +81-298-53-5362
Fax: +81-298-53-5207
E-mail: iwata@kz.tsukuba.ac.jp

University of Washington
Department of Electrical Engineering
Seattle, WA 98195
POC: Dr. Blake Hannaford
Telephone: 206-543-2197
Fax: 206-543-3842
E-mail: blake@isdl.ee.washington.edu

A.2.6 Full Body Motion Interfaces

Computer Graphics Systems Development
Corporation
2483 Old Middlefield Way #140
Mountain View, CA 94043-2330
POC: Dr. Roy Latham
Telephone: 415-903-4922
Fax: 415-967-5252
E-mail: rlatham@cgsd.com

Cybernet Systems Corporation
1919 Green Road, Suite B-101
Ann Arbor, MI 48105
Telephone: 313-668-2567
Fax: 313-668-8780

Flogiston Corporation
Austin, TX
POC: Brian Park
Telephone: 512-894-0562
Fax: 512-894-0562
E-mail: floman@bga.com

Sarcos Research Corporation
360 Wakara Way
Salt Lake City, UT 84108
POC: General Peter Kind
Telephone: 801-581-0155
Fax: 801-581-1151
E-mail: P.Kind@sarcos.com

Denne Developments, Ltd.
Unit 4, Cedar Park
Cobham Road
Ferndown Industrial Estate
Wimborne, Dorset BH21 7SB, UK
POC: Mr. Vatcher
Telephone: +44-1-202-861661
Fax: +44-2-793-432625

University College London
Department of Computer Science
Gower Street
London WC1B 6BT, UK
POC: Dr. Mel Slater
Telephone: +44-171-975-5555 ext. 5242
Fax: +44-181-980-6533

University of Tsukuba
Institute of Engineering Mechanics
POC: Dr. Hiro Iwata
Telephone: +81-298-53-5362
+81 298 53 5362
E-mail: iwata@kz.tsukuba.ac.jp

Systran Corporation
4126 Linden Avenue
Dayton, OH 45432-3068
Telephone: 513-252-5601
Fax: 513-258-2729

Visual Systems Laboratory
Institute for Simulation and Training
POC: Dr. Jim Parsons
(407) 658-5058
Telephone: 407-658-5058
E-mail: parsons@vsl.ist.ucf.edu

A.2.7 Olfactory Interfaces

Artificial Reality Corporation
P.O. Box 786
Vernon, CT 06066
POC: Dr. Myron Krueger
Telephone: 203-871-9445
Fax: 203-871-9445
E-mail: myronkru@aol.com

Marketing Aromatics, Ltd.
#10 Russell Gardens Mews
London, W14 8EU
London, England
POC: Simon Dyer
Telephone: +44-171-602-5632
Fax: +44-171-602-5632er
E-mail: 100572.1472.compuserve.com