Pixels are Good

Mark Bolas, Ian McDowall, George Williams, Dan Corr, Julien Berta[a]

ABSTRACT:

This paper presents three methods for classifying and qualifying virtual and immersive environments. The first is to plot modes of use against environment types. The second is to create a matrix analyzing display type against interaction methodology. The third is to analyze the system as if it created a volume of three-dimensional pixels and determine if the quality of the created pixel volume is appropriate for the given application and use.

Keywords: Immersive Environment, Virtual Environment, Virtual Model, Immersive Display, Head-Coupled Display, Stereoscopic, Presence, Pixel Dust, Virtual Pixel.

1. IMMERSIVE SYSTEMS

Immersive systems are quickly being adopted by many different industries with each application demanding a different type of system configuration. In general, most configurations are composed of two subsystems - an *environment synthesizer* coupled with *environment transducers.*

The environment synthesizer is typically a graphical workstation that includes an audio synthesizer. There is a large family of environment transducers including hand and head trackers, glove and prop based input devices, and head-coupled or head-mounted displays.

While this paper focuses on immersive environment systems that are predominately visual, systems which have both an environment transducer and environment synthesizer need not have a visual component. For example, surgical simulation systems exist for dexterity training which consist solely of a haptic rendering engine and a force feedback device[1].

2. DISPLAYS AND INTERACTION

In general, the higher resolution display with the larger image area or field-of-view, is superior for a given application. With head-coupled displays, it is interesting to observe that tracking can be used to increase apparent screen size and resolution. For example, a 1000 by 1000 pixel display with a 40 degree field-of-view can utilize tracking to change displayed information based upon where the user is looking. In this fashion, the display can emulate a 360 degree spherical display system surrounding the user with an effective resolution of (360 degrees / 40 degrees * 1000 pixels) * 2 = 9000 by 9000 pixels. This creates 18,000,000 effective pixels with a system compact enough to work in a standard office cubicle.

When interacting within immersive environments, physical presence is highly beneficial. This is typically achieved by the use of real objects in a virtual environment. Examples include Ken Hinkley's work with Randy Pausch at the University of Virginia on Passive Props[2] and two handed touch glove interaction devices as pioneered by Dan Mapes and the Toy Scouts at the University of Central Florida[3]. Such interaction devices - especially props - can increase the perceived fidelity of an immersive display by allowing the user to interact with the *virtual pixels* in the user's working volume that represent a virtual object.

[a] Correspondence for all authors: Mail: Fakespace Labs, Inc. 241 Polaris Ave., Mountain View, CA 94043; Email:bolas@well.com; Telephone: (650) 688-1940; Facsimile: (650) 688-1949

For example, if a user is holding a clear interaction palette used to track and tactilely represent the surface of a virtual three dimensional menu displayed via a stereoscopic projection system[4,5] , the interaction palette greatly helps to create the illusion of the projected stereo image as being coincident with the interaction palette. The pixels appear to emanate from the surface of the palette – thus greatly increasing the perceived spatial fidelity of the stereoscopic projection system.

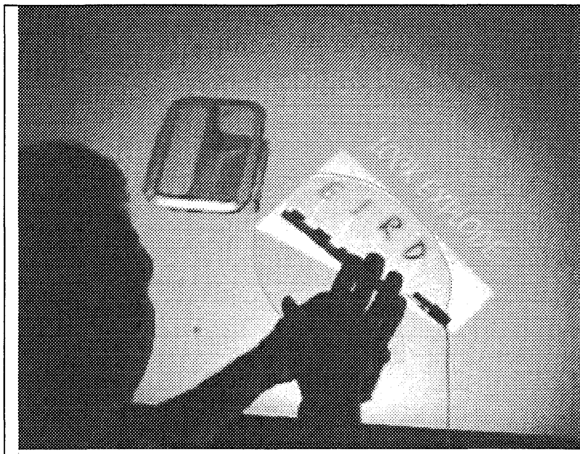

Projected image appears through palette

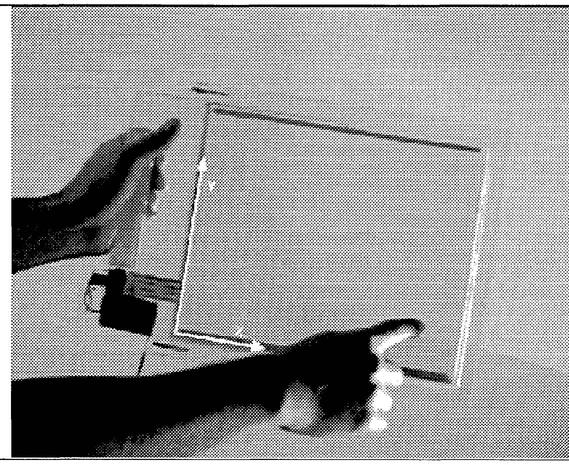

Allows for direct pixel manipulation

3. VIRTUAL PIXELS / PIXEL DUST

The concept of a virtual pixel floating in space before the user and used to represent coincident virtual objects is central to this paper. A way to think of the virtual pixel is to think of all objects within a virtual environment as being comprised of *pixel dust* – imaginary 3D flecks of light floating in space that can be placed anywhere within a virtual environment to represent virtual objects.

The effective size, quantity, shape and spatial characteristics of the granules of pixel dust quickly become interesting limitations to the designer of virtual environments. With this in mind, the description of the granules can be used to qualify different types of virtual environment systems with a uniform set of metrics. It can also be used to test the utility of proposed system configurations for a given application.

For example, it is very difficult to quantify the usefulness of stereoscopy in an immersive environment. Most users will report that a stereoscopic immersive environment 'feels more real' – but this anecdotal data is not useful when trying to justify the increase of a stereoscopic system's procurement cost. Painstaking human performance studies often show decreases in task performance times when stereoscopic displays are compared with monoscopic displays – but the tasks are often only peripherally relevant for alternative applications not included in the original study.

If stereoscopy is thought of in terms of granules of pixel dust, however, its utility is made more obvious. Assume that a stereoscopic display is being considered for a specific application and that the system can display 50 different levels of depth (the expected separation between images for the left and right eye is between 0 to 50 pixels, and a separation of one pixel is noticeable by the user). In this case, thinking of the system as consisting of granules means that any given pixel on the stereoscopic display can be used to represent one of 50 different depth locations. Thus the 1000 by 1000 by 2 display effectively becomes a 1000 by 1000 by 50 display - one million pixels can be placed into 50,000,000 locations.

While this pixel dust analysis points out the gain possible with a stereoscopic display, it can also highlight limitations. For example, it inherently acknowledges a quantification of depth resolution that is often ignored. In addition, it forces the immersive environment designer to choose a display that is the appropriate size for the application's required working volume – there is no reason to locate pixel dust where it cannot be seen or used.

4. PIXELS IN SPACE - OBSERVATIONS

When qualifying systems using metrics based upon granules of pixel dust, it is not necessary to define the elusive *level of immersion*. It is necessary to consider the virtual pixels as if they were a tangible part of the immersive system. Important questions to answer regarding the pixel dust include the following:

- Is it where you need and expect it?
- Is it dense enough?
- Can the user's eyes converge on it?
- Do physical props work with it?
- Is there enough of it?
- Is it where the user can get to it?
- Can you navigate through it?
- What effective working volume does it fill?
- Can the user's eyes accommodate it?
- What artifacts does it introduce?
- Can it be shared simultaneously with other users?
- Can it be seen and presented to a large group?

5. APPLICATION REQUIREMENTS – OBSERVATIONS

Immersive environment systems are used for a broad array of applications. It is important to honestly determine what is critical to a given application, and how the system is expected to be used. To do this, it is helpful to look at the basic *modes of use*, and the basic *environment types*.

There appear to be three modes of using immersive environments. The first is when a single user is analyzing existing data or creating a new design in a concerted and individualized fashion – the system is being used to *contemplate*. The second is when two or more users are working together to analyze data or create a new design and both users are actively engaged with the environment – here the system is being used to *collaborate*. The third is when data or a design is being explained to a group of people with little interaction of the group with the environment – the system is being used to *present*.

Text, *virtual models* and *virtual environments* are three common types of environments. Text – plain old letters and numbers - is very important for many applications and is often ignored. Virtual models are distinct from virtual environments and unless both types are required, a performance gain can be achieved by excluding the other. For example, there is no reason to surround a user with a Cave when the user will only be designing individual small parts. Creation of a matrix similar to the following is useful to see the interdependency between modes of use and environment types for a given application.

Text			
Model			
Environment			
	Contemplate	Collaborate	Present

6. DISPLAY AND INTERACTION MATRIX

Virtual pixels are best created through the combination of a display with interaction devices – not by a display alone. The addition of an interaction device can be used to increase the perceived depth fidelity of the pixel dust, increase the working volume, and so forth. This paper discusses a small cross section of devices and displays with the hope that the underlying analysis can be extended to most immersive systems. The following matrix formalizes the process by highlighting possible combinations – some which will work well, and others that should be avoided.

Boom				
Push				
Bench				
Wall				
Room				
	Gloves	Props	Controllers	Duo

7. EXAMPLES

Projection based room displays, such as the Naval Research Laboratory's Grotto or the CAVE[6,7] are very good at putting pixels where the user expects them. These systems physically build a large work-volume into which they pour four to five screens full of 1,000,000 pixels each. The spaces tend to physically correspond to the environment being represented. As such, interaction is simple - the user turns the body to look at what is of interest. Head-tracking is only used for stereoscopic rendering, with the entire wrap-around environment being rendered at all times. Room displays such as this are computationally expensive and physically expensive as well – all that pixel dust takes up a lot of room to create and see. While excellent for contemplative applications, the current limitation to a single - or at most two - active users makes room systems far less attractive choices for collaborative work and group presentations.

Room displays are typically very good for the visualization of virtual environments in which the user is surrounded by the environment (and thus pixel dust). They are not optimal for the visualization of virtual models[8] or text-based data where virtual pixels go to waste as they are located behind the user or are not dense enough to represent text.

Projection based wall displays typically present 3000 by 1000 pixels with stereo capability on a 20 foot wide surface. They are excellent at the visualization of models that virtually fill the space immediately in front of and behind the screen – typically automotive applications. The ability to go into a non-stereo projection mode allows for the easy projection of text-based data, while the flat and open form-factor allows for group presentations as well as collaborative and contemplative work. Wall type systems are not appropriate however, for visualization of most virtual environments (as opposed to virtual models) as the pixel dust cannot easily surround the user. Interaction tends to be accomplished by the user physically walking around the virtual model, although wand and push type devices are being tested and appear to be useful.

Wall based systems inherently have ambiguous virtual pixels near the edge of the screen due to the reverse occlusion which occurs there. Software can be written to help reduce this effect or a hardware baffle can be used[9,10]. This ambiguity is a good example of a stereoscopic artifact that is recognized when thinking of the system as projecting virtual pixels into space as opposed to merely projecting two sets of time alternating pixels onto a screen.

In order to accommodate the need to surround users with pixel dust - such as with a CAVE, or place the pixel dust into a large virtual model – such as with a wall, systems which can be physically reconfigured to more closely match the desired environment are becoming popular. The RAVE and the MD_Flex are examples[11].

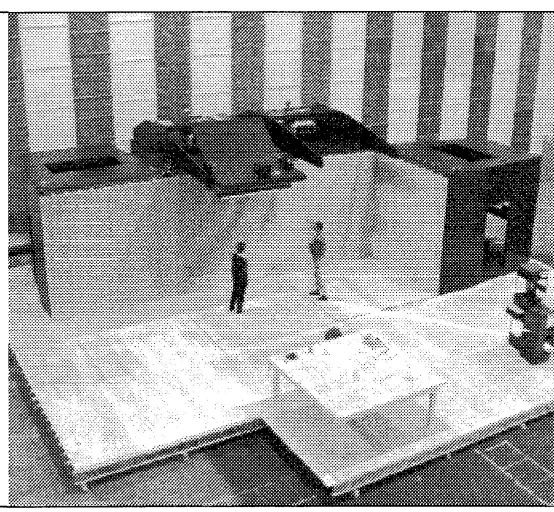

Reconfigurable Cubit/RAVE displays

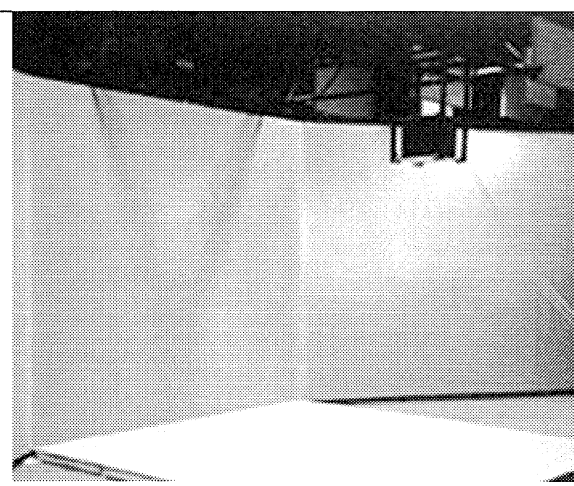

MD_Flex Reconfigurable System

Virtual pixels help explain why workbench systems - and all projection based displays - have a limitation in the usable working volume of the display. While the user's eyes are accommodating on the projection screen, the calculated disparity is typically for a point in space that is in front or behind the screen – for where the virtual pixel is located. The user is focusing on the real pixel emanating from the screen surface, but is trying to triangulate the eyes on the floating virtual pixel. This rivalry between accommodation and convergence places a physiological limit on how far from the screen's surface the virtual pixel can be placed without causing visual discomfort. For medium sized workbench systems, this is typically a foot in front of the screen and three feet behind the screen.

Minibench – Can Reach All Pixels

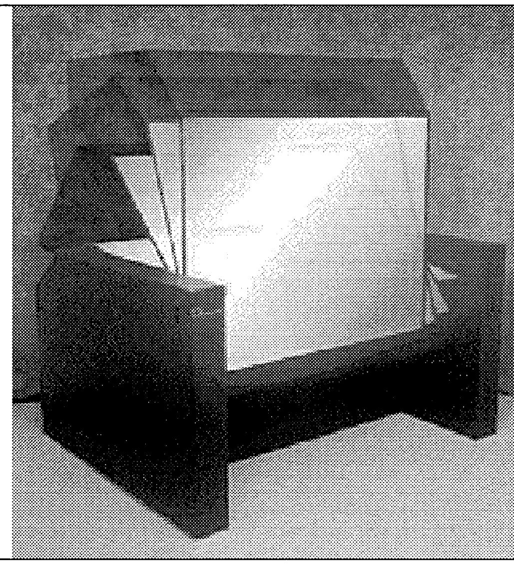

Versabench – Horizontal to Vertical

Workbench systems[12] excel at the presentation of small models due to the tight working volume of their virtual pixels - the high density of pixels can provide stunning resolution as they are concentrated only where needed. This effect is enhanced by the use of direct interaction props such as gloves or palettes. Such props work well on bench systems as the entire volume of pixel dust is within arms' reach of the user. While flat workbench systems work well for a single user doing contemplative work, they fail for use by more than one user - the virtual pixels are in the wrong place for the non-spatially tracked users. The Duo system[13] overcomes this by allowing two sets of mathematically correct virtual pixels to be presented to different users by time sequencing the displayed views. This allows bench systems to be used for collaborative work. For presentations, most workbench systems adjust to provide a vertical display surface that moves the pixels to where they can be seen by a large group.

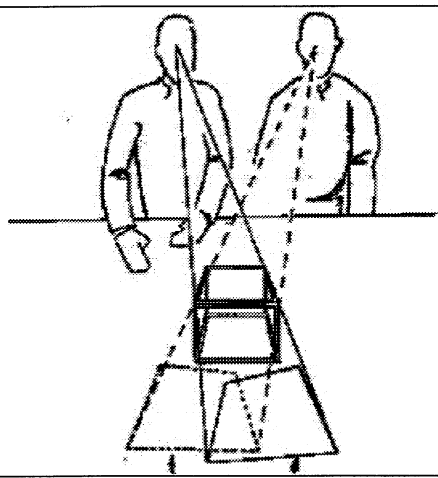

Projected Geometry Shift for Two Users

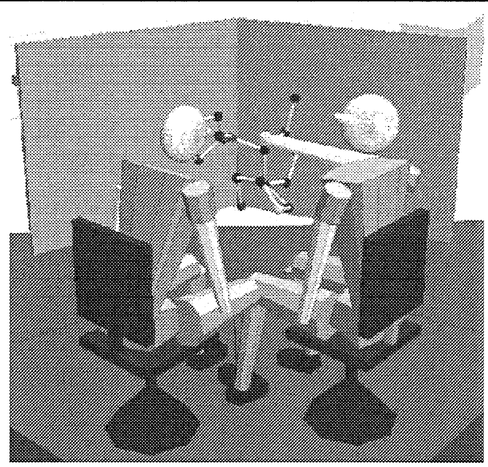

PIT- Two Merged Yet Unique Pixel Volumes

The University of North Carolina's PIT display allows for co-located collaborative work not by time sequencing the pixels but by having two sets of pixel dust overlap in space. Because the user is effectively orthogonal to the other user's screen, each user is given a unique volume of pixel dust. Because each user's pixel dust overlaps the other's, the system ingeniously allows for collaborative interaction[14].

Tracked displays such as the BOOM or a head-mounted display need only display the portion of the virtual environment that the user is looking at. In this fashion, resolutions far exceeding that of projection based displays can be had at a fraction of the cost. Such systems are limited to contemplative work as only a single user can look into a direct view system at a time. By mounting a large flat panel display on a counterbalanced and tracked structure, many of the advantages of the BOOM can be retained while allowing for collaborative viewing between two users. It is interesting to note that even with a monoscopic LCD panel, the ability to physically move the display through space offers the feeling of painting a virtual volume with pixel dust.

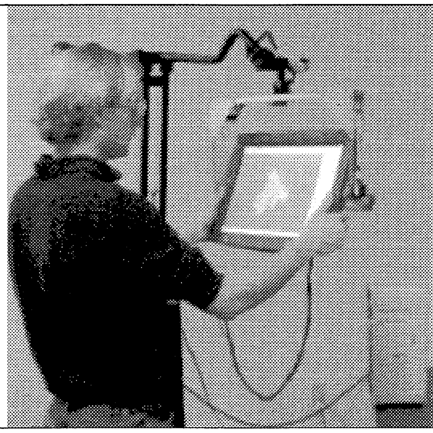	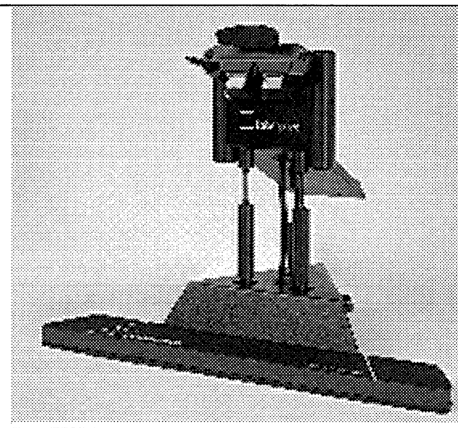
Counterbalanced & Tracked LCD Panel	LCD Based PUSH Display

The Push display tightly integrates a stereoscopic display system with a force-based form of navigation[15,16]. Users can easily navigate around virtual objects as small as a thimble, while simultaneously maintaining navigational control in spaces as large as 100 cubic miles by using their body to push against the display and signal the immersive system to translate or rotate in a corresponding fashion. This effectively allows the user to place a densely filled pixel dust work-volume throughout an extremely large area. Command control and communication applications demand such large work volumes.

Interaction devices such as the WorkPalette allow the user to directly manipulate virtual pixels by placing the finger against the surface of the palette in space and touching the corresponding virtual pixel which is visible through the clear plastic of the palette, yet appears to be located on the palette. Despite the accommodation and convergence rivalry issue, this allows the user to more easily perceive virtual pixels as floating in space, and allows for complex interactions such as streamline placement to be accomplished with simple wrist and hand gestures that appear to be influencing the floating pixel dust.

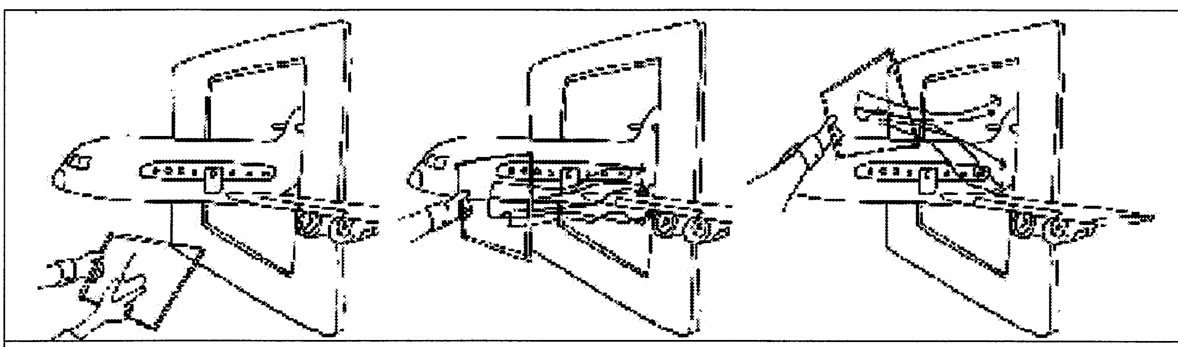

Palette Used to Sketch and Place Streamlines – Used as a Tactile Constraint

8. CONCLUSIONS

Three tools have been presented which can be used to analyze immersive systems. The first is to plot modes of use against environment types. The second is to create a matrix analyzing display type against interaction methodology. The third is to analyze the system as if it created a volume of three dimensional pixels and determine if the quality of the created pixel volume is appropriate for the given application and use.

9. ACKNOWLEDGEMENTS

The authors wish to acknowledge the following individuals and programs: Oliver Riedel, FHG; Steve Bryson, NASA AMES; Russell Taylor, Frederick Brooks, Mary Whitton, Michael Meehan, UNC; Ken Hinckley, Randy Pausch, University of Virginia; Bill Buxton, Alias/Wavefront; Pat Hanrahan, Stanford; Larry Rosenblum, NRL; NAS Contract NAS2-98975; ONR Contract N00014-99-C-0122. MD_Flex image provided by the MechDyne Corporation and the PIT Image provided by the UNC Chapel Hill Department of Computer Science.

10. REFERENCES

1. Dr. Kenneth Salisbury, "Haptics: The Technology of Touch", SensAble Technologies Website: http://www.sensable.com/community/haptwhpp.htm, Nov. 10, 1995 updated Jan 7, 2000.
2. J. Goble, K. Hinckley, R. Pausch, J. Snell, and N. Kassel. Two-Handed Spatial Interface Tools for Neurosurgical Planning. IEEE Computer, 28(7):20-26, 1995.
3. D. Mapes and J. Moshell. A Two-Handed Interface for Object Manipulation in Virtual Environments. Presence, 4(4):403-416, 1995
4. G. Williams, H. Faste, I.E. McDowall, M.T. Bolas. "Physical Presence – Palettes in Virtual Spaces", SPIE Vol. 3639, Jan 1999.
5. G. Williams, I.E. McDowall, M.T. Bolas, "Human Scale Interaction for Virtual Model Displays: A Clear Case for Real Tools", SPIE Vol 3295, Jan 1998.
6. C. Cruz-Neira, D. Sandin, and T. DeFanti. Surround-Screen Projection-Based Virtual Reality: The Design and Implementation of the CAVE. Proceedings of SIGGRAPH'93, pages 135-142, 1993.
7. M. F. Deering. Making Virtual Reality More Real: Experience with Virtual Portal. Graphics Interface'93, pages 219-225, 1993.
8. I.E. McDowall, S.T. Bryson, M.T.Bolas, "New Developments for Virtual Model Displays", SPIE Vol. 3012, Feb. 1997.
9. M.T.Bolas, "Feather-EdgeSystem for the Immersive WorkWall", Fakespace White Paper, Jan 1998
10. Jurriaan D. Mulder and Robert van Liere , "Enhancing Fish Tank VR", IEEE Virtual Reality 2000
11. Kurt Hoffmeister, Mike Hancock, and James Gruening. "Developing World-Class Immersive Environment Facilities". Presented at the 3rd International Immersive Projection Technology Workshop, May 10, 1999.
12. W. Krüger, C. Bohn, B. Fröhlich, H. Schüth, W. Strauss, and G. Wesche. The Responsive Workbench: A Virtual Work Environment. IEEE Computer, 28(7):42-48, 1995.
13. M. Agrawala, A. C. Beers, B. Fröhlich, I. McDowall P. Hanrahan, and M. Bolas, "The Two-User Responsive Workbench: Support for Collaboration Through Individual Views of a Shared Space", Computer Graphics Proceedings, Annual Conference Series SIGGRAPH '97, pages 327-332, 1997.
14. Kevin Arthur, Timothy Preston, Russell M. Taylor II, Frederick P. Brooks, Jr., Mary C. Whitton, and William V. Wright. "Designing and Building the PIT: a Head-Tracked Stereo Workspace for Two Users." Presented at the 2nd International Immersive Projection Technology Workshop (Ames, Iowa, May 11-12, 1998). Also available as UNC Chapel Hill Department of Computer Science Technical Report TR98-015.
15. M.T. Bolas, I.E. McDowall, R.X. Mead, "Immersive Desktop Display with Axial Muscle Navigation and Control", White paper distributed at '95 Symposium on Interactive 3D Graphics, Monterey, CA.
16. R.P. Darken and J.L. Silbert. A Toolset for Navigation in Virtual Environments. In Proceedings of ACM User Interface Software & Technology, pages 157-165, 1993.