Virtual Reality for Visualisation

Claire Knight

Visualisation Research Group,
Centre for Software Maintenance.

Department of Computer Science,
University of Durham,
Durham, DH1 3LE, UK.

E-mail: C.R.Knight@durham.ac.uk
Web: http://vrg.dur.ac.uk/

Computer Science Technical Report 13/98 – October 1998

Abstract

Virtual reality is an effective way (in many situations) for creating 3D visualisations, although the visualisations need to be designed with both 3D and virtual reality in mind to be most useful. Many issues are related to virtual realities, although these are often overlooked by users due to the issues being tightly integrated with the virtual worlds. Even though users can disregard such issues, it is important that designers of visualisations and virtual realities bear them in mind. Some of these issues are presented and discussed to enable designers to have a better idea of the aspects they should consider when creating such systems, with particular reference to visualisation with virtual reality.

1. Introduction

Software visualisation is an important tool in the maintainer's armoury. It is a technique that can, when designed and used effectively, aid him in his quest to understand existing program code. One way to achieve these visualisations is through the use of virtual reality and virtual environments.

Collaborative virtual environments have also benefited from the technological advances that aided visualisation. Whilst such environments can be created in textual form, it is generally accepted today that such environments have a background of graphics. Indeed some of these systems are very close to virtual reality and games environments.

Virtual reality (VR) provides a way of producing 3D visualisations, and allowing the user to interact with and in those visualisations. Unfortunately it is not as simple as translating 2D visualisations into 3D because the one extra dimension greatly affects the way the data can be presented. Not only do new visualisations have to be designed to make effective use of the extra dimension; they also have to be designed with VR in mind. If a system is not designed to make use of the features that VR can provide then the visualisations produced can be no better (or even worse) than existing ones. This then produces a negative view of VR as a tool, which need not be the case.

2. Collaborative Virtual Environments

Collaborative virtual environments (CVEs) is a research area primarily concerned with the construction of multi-user virtual worlds rather than the work processes or interaction that takes place within those worlds. The research covers all facets of virtual world creation from the underlying graphics engines to varying algorithms for creating the best perceptual experience within the worlds. Because of this focus much CVE research can be linked to Virtual Reality (VR) systems and research.

Ellis [Elli94] defines virtual environments as

> *"…interactive, virtual image displays enhanced by special processing and by nonvisual display modalities, such as auditory and haptic, to convince users that they are immersed in a synthetic space."*

but this is not an adequate description for all virtual environments. This definition assumes that the user is wearing, and the system supports, force feedback devices and headsets. Much of the CVE research done to date, and indeed an increasing amount of virtual reality and visualisation work, does not require this level of hardware.

Churchill and Snowdon [Chur98] provide a very good description of CVEs:

> *"As such CVEs provide a potentially infinite, graphically realised digital landscape within which multiple users can interact with each other and with simple or complex data representations."*

CVEs can be found in science fiction and fantasy literature where general on-line worlds are described rather than being for one purpose, as is the case with most current CVE systems. Examples of these "worlds" are Neal Stephenson's [Step92] *Metaverse* and William Gibson's [Gibs84] *Cyberspace* or *Matrix*.

Sandor et al. [Sand97] define CVEs as

> *"…, where Virtual Reality (VR) technology is used to support cooperative applications such as virtual conferencing or collaborative information visualization and retrieval."*

2.1 Embodiment and Awareness

User representation (embodiment) within both CVEs and general virtual reality systems is still very much an area of research but a general name for these representations is *avatars*. This term is now universally accepted as a description of a person's representation in a virtual environment.

The sophistication of the avatars in terms of the graphics used to display them and the control the user has over them can affect the amount of interaction that goes on in a virtual environment. If avatars are able to show different facial expressions and gestures in a realistic manner then these will help the interactions in the virtual world to be like those experienced in the real world.

Benford et al. [Benf97b] did some experiments with users inside a virtual environment watching a performance by others users in that same environment, all of which was visible to users outside of the environment. This experiment showed that not only was the aesthetic appearance of the avatars important but also that the symbolic representation needed to be considered.

Awareness can relate not only to the current task but also to any other activity that the user is involved in. It may be that a collaborator is away from the virtual environment for some short-term reason, such as (in today's technology at least) sending an e-mail. He may not bother to leave the virtual environment but the other users of the system need to be aware that he is not concentrating on that environment. There is also the case when a collaborator has left the environment but a colleague wishes to contact him the next time that they are both in the world. Being able to leave a message to this effect within the environment can prove invaluable.

These two facets of awareness cover only a user not directly participating in the environment but there are other situations where knowing what other users are doing can be useful. Knowing that a collaborator is reading, talking to someone else or focusing his attention in a particular place are all useful pieces of information for another user. Bentley et al. [Bent93] write:

> *"In contrast, cooperative applications need to provide users with an awareness of the activities of others to support and encourages cooperation to take place."*

which for CVEs could be considered to be more important than for other computer collaborative mechanisms.

In the experiment conducted by Benford et al. [Benf97b], the level of awareness of participants was important. There was a central *region,* which was the stage. The actions and audio from the performers could be seen and heard by all users of the virtual environment whilst there were smaller local regions surrounding non-performing users of the system. These smaller regions allowed those in them to see and hear others in the same region and the rest of the virtual environment but those outside that region could see it only as a coloured cone. The awareness is handled through the user of third party objects, which are independent objects in a virtual world that can perform awareness adaptations and aggregations.

Sandor et al. [Sand97] have developed a different method of awareness known as Aether, and is considered to be part of the system architecture of the virtual environment. The awareness engine (Aether) is then able to provide applications with information about the users current (and past) activities. A network of objects interconnected with direct relations is maintained which is used for the awareness computations. The objects can be any entity that can exist in the environment.

Aura, nimbus and focus are also important when considering awareness in virtual environments. Aura is the potential for collaboration between objects and can be seen as the scope or area of interest. It is a sub-space that effectively bounds each object. Nimbus and focus allow the objects to control the level of awareness by and of themselves. These are also sub-spaces bounding an object. The focus provides a measure of how aware an object is of anything else whilst the nimbus is used by other objects as a measure of their awareness of that object.

2.2 How Views Affect Multiple Users

Early virtual reality systems (and in some domains such as the medical field it is continuing) were only capable of presenting the same view to all users. The applications were designed to allow one user to effectively control a tour or presentation of information. This is not appropriate for virtual environments that need to support multiple users in a collaborative nature. There is a need to be able to provide tailorable subjective views to different users. Churchill and Snowdon [Chur98] also make a point about the business value of the data being viewed and manipulated in the virtual environment.

> *"There is also a broader social context within which visualisations may need to vary; if the data being visualised is commercially sensitive, then viewing certain pieces of information might need to be restricted to a subset of the user population. If the virtual environment does not support subjective views then the users are forced to agree on a common (possibly non-optimal) visualisation style."*

Bentley et al. [Bent93] summarise the three levels of sharing in current co-operative systems which correspond to different levels of user interface coupling. They define this coupling as

> *"Different forms of cooperative work require varying levels of awareness between users and place different demands on the strength of sharing required. The extent to which multi-user interfaces support this sharing through the propagation of activities is termed* interface coupling*; the greater the level of awareness between users the closer the interface coupling."*

The three levels of sharing are:

- Presentation level sharing (Tight coupling)
  Each user is presented with the same display of the environment
- View level sharing (Medium coupling)
  Each user is viewing the same area of the environment but the actual individual displays may differ.
- Object level sharing (Loose coupling)
  Each user may be viewing different areas of the environment in different ways.

In many situations the object level sharing is best for the virtual environment despite it being classed as louse coupling, hence a lower level of awareness. Work on awareness engines since this was written have shown ways in which individuals can have the views that best support their current work, but still allowing effective communication and awareness between users.

3. Virtual Reality

Virtual Reality (VR) has become a hot topic in recent years. The increasing power available in today's desktop computers and consequently their ability to display better graphics has pushed this field into the public eye. The more advanced industrial and research systems have very sophisticated machinery available to them including spatial monitoring of users and haptic devices for use with the VR systems.

3.1 What is Virtual Reality?

A good definition of Virtual Reality (VR) is given by Isdale, [Isda93] (referenced from Aukstakalnis et al. [Auks92]):

> *"Virtual Reality is a way for humans to visualize, manipulate and interact with computers and extremely complex data."*

This definition puts the emphasis in computer representations of some form, along with the ability to be able to interact and manipulate the visualisations provided in some way. Some people view VR as purely the special headsets and control pads/gloves and nothing more. Others see VR as much more, such as reading a book where the mind creates the virtual environment. These two views are extremes and the definition provided above takes a middle view and the one which best encompasses the computer science perspective when VR is taken as a tool rather than a technology to be investigated.

There are several types of VR systems around. These range from the simplest display on a normal computer monitor to fully immersive systems. The systems that make use of computer monitors as standard are known as “Window on a World” systems (WOWs). Many of the computer games available today fall into this category – especially the first person viewing systems such as Wolfenstein, Doom, Quake, Descent and many of the desktop flight simulators. Various other games make use of 3D (especially the sporting simulators such as the FIFA range of games made by EA Sports) but generally the *camera view* is not one which lends itself to giving the user a perception of VR. It can be argued that it is just a different view provided by a WOW, a more passive viewpoint where there is control over the contents of the window but the camera views and direction of looking are controlled elsewhere.

Video mapping is a technique that is used to put an image of the user into the world and the user can then view his own interaction on a standard monitor using the WOW method.

The total VR system is one that immerses the user (viewpoint) within the virtual world. Special technology is needed to be able to achieve this. These systems use head mounted displays that contain all the visual and auditory information necessary for effectively interacting in the virtual world. A variation is when a room is converted into a complete visual environment using several large projections of the virtual world the user is immersed in. There also needs to be appropriate tracking mechanisms and a means of interaction with the user.

There are ways to combine VR displays and worlds with information visualisation of a more conventional nature. In this way even more information can be presented to the user. A simple example is having a map of the virtual world displayed in one corner (or easily accessible from a movement/trigger). Base information is then combined with the virtual experience.

The equipment for immersive VR systems is getting more sophisticated as the available hardware improves. Instead of just head mounted displays and sensory gloves there are now full body sensor suits making use of position sensors allowing more sophisticated worlds to be created and experienced. Another use of the position tracking of humans is to create fluid, realistic animation in computer games and simulations. EA Sports and Gremlin Interactive have both used this technique with their football games (FIFA 97 and Actua Soccer respectively).

Immersive VR relies on stereo vision. The brain constantly generates information such as depth perception based in its fusing of the images from each eye. VR technologies play on this by causing the brain to do the hard work. Two images, differing slightly, are presented to the user by means of two monitors, or in front of each eye in a head-mounted display. The brain, when the images are processed quickly enough can cause the perception of depth. Generally a refresh rate of greater than 60hz is required so that there is no flickering for the user.

Rendering in a VR system needs to exceed 20 frames a second since this is the minimum level that the brain will take a series of still images and perceive smooth animation. The experience of reality is also enhanced by the use of audio. 3D audio, implemented correctly and accounting for the problems of the brain trying to work out the placement of the sound in the mind, can make the virtual world more like what is known as reality. 3D audio can also add to the perceptual experience and may add to the realism. There have not been enough studies or examples of work in this area to conclude one way or another.

Recently some studies have been done that claim the closeness of head mounted displays cause unnecessary stress and vision strain in the wearer. The claim is that it is the closeness of the images, and forcing the brain not to focus on the close object but to generate depth perception cause the problems.

3.2 Basic Perception

Perception is the observation of the surrounding world. It can be used in part explaining why things appear as they do. In principle, science is capable of explaining the world – biologically with vision and wavelengths of light from physics. Perception needs to be understood to facilitate the creation of an environment. It can also be of benefit in understanding observational errors, the senses upon which perception depends and for creating a machine to simulate human behaviour in some way.

Perception is never exactly in accord with the reality in which that perception is made. The human brain is very good at missing details, distorting them, or even making the eye see what is not there. Some of the missing detail is necessary to cut out the complexity. Surfaces of some objects are recognised as shiny but the brain (unless being specifically focused on the reflection) does nothing more than provide a hazy mirrored view. This type of missing detail is not in error. Optical illusions play on this idea of confusing the brain and exploit defects. It is these facts that can allow virtual environments to be created with some success.

By knowing and understanding perception in general, rather than in a particular case, things that are perceived easily or that trick the brain can be used to good effect. In addition if something is noticed by our senses then there is still no guarantee that it will be perceived. Again this is exploited with optical illusions. Information on the visual ability of humans and their use of vision for recognition can be found in the writings of Findlay and Newell [Find95]. They cover several approaches to recognition and relate this to vision. They also acknowledge the fact that the human visual system (be it eye, brain or a combination of these) is very good at recognising things even from crude (basic) images.

Perception of life is an individual phenomenon. With exactly the same stimuli two people may respond completely differently. Everything that has gone before in a person's life will affect their perception of a situation. This indicates that perception is more than just vision or spatial awareness. Existing knowledge stored in the brain can also be influential in determining the perception of a given event.

Pettifer and West [Pett97b] discuss perception (and philosophy) in relation to the use of metaphysical models in virtual worlds. They make the point

> *"Various commentaries on the distinctions between the 'objective universe' and our experience or perception of it have been made throughout human history (indeed, the very question of whether or not there is an 'objective' existence outside our perception of it is one that underpins many philosophical works)."*

They then go on to say that Immanuel Kant (German Philosopher, 1724-1804) are most appropriate when discussing perception in relation to virtual reality and virtual environments. This is because he made the distinction between "the thing itself" and "the thing for me". The first is the objective nature of something which humans cannot directly experience whilst the second is the subjective experience, something humans are able to perceive.

Perception is composed of various "inputs". All of the senses contribute to an overall perception of some event. Seeing and hearing contribute to distance perception, touch (skin) deals with closeness to objects and the position/motion of muscles and joints (the skeletal and muscular system) provide motion and position information. Additionally if something does not affect one or more of these sensory inputs then nothing will be perceived.

Learning to perceive can and does happen. A person's face becomes recognisable as a friend or a colleague when the brain associates a perception of that person with knowledge that is knows whom that person is. This in itself is more than just recognition. The brain (and eyes) are then capable of looking for this person in a crowd using features from the face. This means that the brain is capable of complex perception. There is the knowledge of being in or near a crowd and there is the scanning of faces within the crowd trying to locate a familiar face.

There are many other examples where learning to perceive something has a beneficial effect. Learning a new language (both written and spoken) requires perception of both images on a page and sounds in to the ears. All learning provides more perceptual knowledge. Reading diagrams from a particular field can be virtually impossible without knowing what particular symbols means. Once there is this knowledge and using several diagrams the brain learns to perceive the mass of printed information as something meaningful.

Gardner [Gard93] summarises this perceptual learning through experience as:

> *"These perceptual associations are not so strange. People pick up a lot of rules as they go through life. Not necessarily hard and fast rules; people use a kind of fuzzy logic. When we see something, we relate it to other things."*

In deciding to move in a particular path there is very often the ulteria motive of wanting to more from one location to another, not just to move along that route. It is the perception of the space/world that allows the decision on which path to take. Generally the physical world properties are not compelling in their own right, other than to avoid obstacles when moving from one place to another.

Friedhoff and Benzon [Fri91] write (p12):

> *"The tendency to impose form, whether on the surface irregularities of a cave, or in inkblots, clouds or shadows, is suggestive of an organising function of vision – an organising tendency so strong that random shapes can trigger the perception of vivid illusions."*

This information is something that can be exploited, both by visualisation and those wanting to learn more about the human brain. In better understanding the way the brain and eye perceive colours and images then the better the visualisations that can be created.

Aesthetics is an important factor to consider in relation to perception. Aesthetics very often determines how long something is viewed for either because it is *aesthetically pleasing* or because it is appalling! It is perception that lets the viewer know what they think is aesthetic. Aesthetics also relates back to the philosophical ideas of Kant (included in [Pett97b]) because it is a subjective judgement. By Kant's definition this means humans can directly perceive it, rather than anything objective, which he says, is impossible to directly experience.

In presenting spatial information, i.e. visually, there is the benefit that the human perception skills can be used for part of the comprehension. This moves some of the comprehension load away from the conscious cognitive processing.

3.3 Stereoscopy

When humans (and many more living creatures on this planet) use their visual systems the image that the brain thinks it is seeing is based on two slightly different views; one from each eye. This use of two images to provide the complete picture is where stereoscopy is relevant.

Stereograms are simple, fun examples of this. To view them there is a need to defocus the eyes because the two different pictures (that create the depth) are encoded into one picture. If the eyes focus as normal a flat pattern is seen. When the eyes are "correctly" focused a three dimensional image comes out of (and recedes into) the original picture.

The earliest examples of stereograms were stereo pairs. These consisted of two images (photographs or paintings) side by side, with very slight differences. By looking at each picture with one eye the brain was able to fuse together the two images to create depth perception.

Two image versions of these (using photographs and paintings) have been around for about a century. It is only with the advent of better computer graphics techniques that the single images containing a 3D image have become popular and widely available. One of the problems with this form of 3D imaging is the inability to make use of the third dimension to aid in visualisation.

Single Image Random Dot Stereograms are a subset of stereograms and are sometimes known as autostereograms. These started with Random Dot Stereograms that are based on the apparently random placement of dots. This form of encoding three dimensions on a two dimensional surface was developed in 1959 by Dr. Julesz. Later Tyler developed an algorithm for single picture images, now known as SIRDS. Tyler studied Psychology at Leicester and Aston universities and received his Ph.D. from Keele. In 1979 he presented the first SIRD.

Stereograms make use of the brain fusing two images together to present a three dimensional image in a 2D printed page. These became very popular in the early 1990s and there are now several books available containing these images. Simplified versions of these are Single Image Random Dot Stereograms (SIRDS) that are simply dots on a page. The images can then be seen by allowing the eye and brain to fuse together the two images hidden in the dots. The only problem with these methods is that the 3D image only has the colours of the 2D image. Depth perception is created but there is no way of encoding colour in the separated images. The principle of stereograms lies in presenting the viewer with two slightly different images, one for each eye.

There is a strong relationship between stereography and virtual reality. Each uses a different picture for each eye and the brain is fooled into seeing a three dimensional image.

Ware [Ware95] writes that the human stereo processing mechanism is highly flexible and stereopsis provides only local additional depth information. This paper presents an algorithm that allows for the dynamic adjustment of stereo displays, one advantage of which is reduced eyestrain for the user.

Stereoscopic depth cues are provided by differences between the two images (usually each views a slightly different image). If the differences between the two images become too much then a problem known as *diplopia* occurs. This is when the images can no longer be fused into one image and there is an effect of parts of the image being doubled. This does not completely destroy depth perception and some effective judgements can still be made.

Motion Parallax is the depth effect created when the visual processing part of the brain is being given information that it is unable to use for stereopsis. This is known as a visual flow field for the simple reason that the effect occurs when moving through an environment where stereopsis is ineffective because of the range of objects being viewed.

Depth cues provide most of the human understanding of 3D space. Motion parallax, occlusion (closer objects obscuring farther ones) and perspective are all used by the brain to provide a perception of depth. Stereopsis provides additional information and is not the main source of depth information.

The amount of difference between images (the disparity) is important. Each human has their own level of comfortable disparity which leads to the conclusion that systems making use of such techniques as stereoscopy should allow some configuration for the comfort of the user.

3.4 Perceptual Vision

The above detail on stereoscopy has provided information on some of the effects of vision, but this is really only of importance for head-mounted display type VR systems. Many more VR systems use standard monitors to display data on, so the effects of vision on the users' perception of the VR world are important.

Most VR worlds are built based on Euclidean Geometry [Cast97], which is based on five postulates. It was later proved that the fifth postulate, which Euclid himself was unhappy with (from a mathematical viewpoint), could be disproved hence creating two new non-Euclidean geometries. One geometry is that of being on the outside of a sphere (which is actually closer to our reality than Euclidean Geometry) and the other the inverse; that of being on the inside of a sphere. Despite reality not being based on Euclidean geometry it is accepted for use in architecture and VR worlds because our limited perception in reality is imperceptible close to it.

There are two relatively modern theories of perception; Gestalt psychology and Ecological optics. These try and explain the way that an environment is seen and comprehended at a perceptual level.

3.4.1 Gestalt Psychology and Ecological Optics

Gestalt psychology has five basic laws of perceptual organisation. These are:

- Proximity
- Similarity
- Good continuation
- Closure
- Common fate

The last four of these are grouped under Prägnanz principles and are concerned with figure perception. Gestalt is the concept that perceptually the whole is greater than the sum of the parts. The relationships between objects are more important than details on those objects. Every book, paper or web site seems to classify the laws or principles in different ways but the above information was provided by Gardner [Gard93].

Cloze testing came from the idea of closure. It refers to the human ability (and indeed inclination) to complete things even if parts are missing. The application of closure (and cloze testing) has been widely applied in the field of text comprehension. Some of this work has been extended to apply to program comprehension and Davis [Davi95] documents a guessing measure of program comprehension which uses some of these theories.

Gestalt theories were developed in the early part of this century (20th). Following this an even newer form of perceptual psychology emerged. This is known as ecological optics, although sometimes it is referred to as Gibsonian theory (after James J. Gibson who created it). Ecological optics takes the view that the eye detects environment invariants, which are things that hold true over a long period of time and the visual system has evolved to over millions of years. Gardner [Gard93] relates these perceptual invariants to VR systems:

> *"To a large extent, the use of perceptual invariants can be programmed into the system. Textures, texture gradients, linear perspective, aerial perspective, motion parallax, … can be part of the graphics and animation system."*

Over time Gibson moved the emphasis from invariants to affordances. An affordance can be seen as a possible action, and Gibson maintained that these affordances could be perceived directly rather than through stimulus identification. Findlay and Newell [Find95] discuss aspects of ecological optics and also the criticisms that have been directed towards the theories, especially since the more recent changes in thought regarding human vision.

Card et al. [Card91] applied the affordance concept of *active perception* in their work on The Information Visualizer. They make use of both animation and three-dimensional graphics to enable easier information retrieval of database texts.

> *"Thus interactive animation and 3D perspective graphics both allow us to apply Gibson's active perception tenants and to pack the space more densely with information that would otherwise be possible. By manipulating objects or moving in space, the user can disambiguate images, reveal hidden information, or zoom in for detail – rapidly accessing information."*

All of these features can be seen to be desirable from a visualisation standpoint, and ecological optics provides a psychological basis for their use.

3.5 Immersion

Immersion was mentioned in the earlier section *3.1 What is Virtual Reality?* The sentence

> *"The total VR system is one that immerses the user (viewpoint) within the virtual world."*

has viewpoint in brackets because the use of the term has come to mean several things in VR. There are two main uses of the word immerse in VR literature. The first is that to be immersive a system has to have a head mounted display (HMD), motion tracking, stereoscopic display and associated paraphernalia. The second, which is now starting to be used more, is that immersive describes the perceptual experience the user has of the virtual environment. This second definition does not rely on specific pieces of technology and therefore the term can be used to describe both traditionally immersive systems and what have been called window on a world (WOW) systems.

Hodges et al. [Hodg95] suggest that "true" immersion is necessary to distinguish VR from interactive graphics and write

> *"User immersion in a synthetic environment distinctively characterizes virtual reality (VR) as different from interactive computer graphics or multimedia. In fact, the* sense of presence *in a virtual world elicited by immersive VR technology indicates that VR applications may differ fundamentally from those commonly associated with graphics and multimedia systems."*

In writing this they are saying that to experience immersive VR, extra hardware technology such as head mounted displays are required.

Bolas [Bola94] also supports this view. He writes

> *"The powerful experience of immersion in 3D visualizations – something "stereopaths" have always known – is now opening up to computer users everywhere. Nowhere is this truer than in the field of virtual reality."*

This assumes the use of stereoscopic projection of some kind (be it projectors, double monitors or head mounted displays) to present the virtual world to the user. Ellis [Elli94] goes one step further and includes other devices, such as haptic interaction tools, in his concept of VR immersion.

> *"In fact, we can define virtual environments as interactive, virtual image displays enhanced by special processing and by nonvisual display modalities, such as auditory and haptic, to convince users that they are immersed in a synthetic space."*

This definition is still reliant on the technology used by the VR system and does not consider the perceptual experience the user is able to obtain.

Gardner [Gard93] makes a distinction between *immersion* and *inclusion* in The Creator's Toolbox, when relating virtual environments to aspects of perception. Whilst he does not provide any evidence for whether the hardware is included in his definition, he does not explicitly say it is required.

> *"Currently, most virtual reality interfaces represent only the hands, if anything, for visual representation of self. To the best of my knowledge, no one has represented the user's feet for calibrating ground textures or even the user's own nose (arguably the most viewed object in a person's visual experience) for self-perception in a virtual reality interface. This distinguishes* immersion *from* inclusion*, which, respectively, give the feelings of being surrounded by the other world and being part of its environment."*

In talking about inclusion, Gardner is touching on what is known as the concept of *degree of presence.* It has also been argued that the degree of presence felt by users is dependent on the immersive level of a VR system. This may well be true, but again some people believe that a user can only feel fully present in a virtual environment if they are using headsets, 3D mice and wands. There is also an associated *degree of presence problem*. This is where users who are in the virtual reality have "virtual out of body experiences". This occurs when the user concentrates on something in reality (at least the nearest there is to our knowledge!). The problem occurs when other users are not aware that they have changed focus and try to communicate and interact with them in the virtual reality. This problem has been documented by several authors, including the work done by Benford et al. [Benf96] in trials with the MASSIVE system.

An alternative view of immersion is documented by Machover and Tice [Mach94]. Their view is that the equipment used does not define a system as being VR (or not).

> *"VR is unique in its emphasis on the experience of the human participant. VR focuses the user's attention on the experience whilst suspending disbelief about the method of creating it. We feel that neither the devices used nor the level of interactiveness or fidelity determine whether a system is "VR"."*

Bowers et al. [Bowe96] are also of the opinion that the term immersion needs to be redefined. Early on in the paper they write

> *"In the conclusion of our paper, we examine the implications of our work for the design of CVEs and suggest – in line with our empirical studies – a redefinition of a central concept much used in VR research: immersion."*

This comment is supported by a discussion later in the paper. The authors suggest that the reader may object to the research and its results due to the consideration of purely desktop VR systems and (in the old meaning of the word) non-immersive VR systems. As part of the defence of this they write

> *"Our* second *response to the objection that we have confined ourselves to non-immersive desktop virtual worlds is to urge a reconsideration of what is meant by 'immersive' in such contexts. We would suggest the utility of understanding 'immersion',* not *in terms of different technical arrangements (e.g. HMD versus screen-based presentation),* nor *in narrowly defined perceptual or psychophysical terms (e.g. is the virtual world all that one can see/hear? To what extent does the virtual world take account of human perceptual systems?),* but *as a practical accomplishment brought off through the work done in giving social activity an orderliness in the virtual world. As such immersion would need to be understood as involving* relations *between one's real-world activities, one's virtual world activities, one's ability to display one's activities in either world to others, relations which are ongoingly achieved, maintained and (if necessary) repaired."*

This is a view taking into account the human factors of virtual reality systems, and realistically, the situations in which VR in the workplace can be effectively used. The use of expensive, specialist, and possibly bulky, equipment is not moving the benefits of VR any closer to the workplace (or to the home). It is also possible to create perceptually immersive environments without the use of such equipment.

To avoid confusion, anywhere in this document where immersive may have been used according to the second definition the term *inside* will be used instead. This is done so that, hopefully, no confusion is caused, and it avoids changing the most common meaning of the term immersion.

3.6 Metaphors

A metaphor is where a word or phrase (or in terms of visualisation, a graphical representation of that word or phrase) is used in place of another. This tends to suggest some form of analogy between the two concepts, although this may be at a higher level of abstraction than individual words or phrases. From a VR perspective the metaphors act as a mapping from the concepts required in the virtual world to their graphical representation. This need was identified by Levialdi et al. [Levi95] in the construction of their database visualisation system.

> *"Using VR visualization techniques to represent the results of queries implies the definition of a mapping, or* metaphor*, among the objects of the database and the objects of some virtual world."*

According to Benford et al. [Benf96] the use of natural metaphors can aid the usability of virtual environments.

> *"… an attempt to exploit people's natural understanding of the physical world, including spatial factors in perception and navigation, as well as general familiarity with common spatial environments…"*

Fitzpatrick et al. [Fitz96] also apply the spatial metaphor to the level of social interaction possible within the virtual world representation of the metaphor.

> *"Even though space is an intuitive, familiar metaphor to work with, there can be a more encompassing meaning of space in the virtual, independent of graphical and VR depictions, that is driven by social world needs and the needs of individuals participating in multiple social worlds."*

Pettifer and West [Pett97a] suggest that the potential power of virtual reality comes from the strength of its metaphor, and the fact that it is closer to natural interaction than many other forms of computer system. They also identify the benefits of natural metaphors, and making use of perceptual and spatial skills learnt and used in the real world in the virtual environment.

> *"A three-dimensional world metaphor has much more scope for direct human/computer interaction than the two-dimensional desktop because it engages in us those perceptual and spatial faculties that allow us to comprehend our surroundings and to process effortlessly the vast amounts of information that are presented to our senses second by second. It is the potential to directly engage these faculties that is the defining characteristic of virtual reality. As the*

immersive environment is far richer than the desktop, the metaphors for interaction assume a far greater significance. ... The role and management of metaphors for the virtual environment therefore assumes key significance."

It is obvious from the above that the design of the metaphor used in the virtual environment can play a large part in the usability of that system, both in terms of human computer interaction, and in terms of enabling the user to carry out the required tasks. What is also of benefit is that in using three-dimensional environments some of the cognitive processing needed for navigation and visual interpretation can be shifted to the sub-conscious as these are activities that are carried out daily with no real thought.

3.7 Spatial Orientation and Navigation

If the VR environment is a representation of the spatial world that we already know then there is a need to model orientation and navigation features found in the real world. In any spatial setting some form of base orientation needs to be found which can then be used for navigation and re-orientation as movement occurs. Hemmje et al. [Hemm94] relate this to their database visualisation work although what they write is readily extendible to all spatial visualisations.

> *"It is necessary to* move*, i.e. change position in the context space and* explore *information visible from each point of view. It is important to achieve an* orientation*, i.e. to determine the relation between a* current point of view *(e.g. from an information item) and the whole of an information space."*

Many authors document the problems of getting lost in "cyberspace" when dealing with spatial virtual environments. Ingram and Benford [Ingr95] write

> *"More recent experiences with virtual reality suggest that users will also suffer from the commonly experienced "lost in hyperspace" problem when trying to navigate virtual environments."*

They relate the orientation and navigation processes to the cognitive map the user has of the environment. Cognitive maps can be one of two sorts. Linear maps are based on movement through the space and the observations made during that movement. Spatial maps do not require movement through the space. Generally linear maps are the first created of an environment, and over time the map may evolve to being a spatial map. Exploration rather than guidance through an environment encourages the development of a spatial map. Their research has focused on providing ways to ease the navigation (and orientation) problems that occur in VR.

Pettifer and West [Pett97a] also relate the problem to the systems and metaphors in use today.

> *"Loosing a cursor on the desktop is one thing, loosing yourself in cyberspace is quite another."*

Three-dimensional worlds are potentially infinite whereas desktops are of generally finite space even if current implementations are able to cover several screens.

Hubbold et al. [Hubb93] discuss design issues that are important to consider for virtual reality systems and cover orientation when discussing perceptual consistency.

> *"More important is the creation of an environment in which the user remains comfortable and well oriented."*

Pettifer and West [Pett97b] also comment on the construction of virtual environments, and that the aim must be to construct these environments so that they correspond with human perceptual requirements. Backing up these comments made by the above authors, Pesce [Pesc93] asserts

> *"The first perogative in the engineering of a holosthetic environment is: design to avoid disorientation. Disorientation represents a step towards the amputation of the self, and necessarily precedes the dislocation of self that concludes in holosthetic psychosis."*

Another aspect of perceptual orientation, often missed, is that of causality. It provides a continuity of experience in "reality" so by providing such continuity in virtual realities allows natural comprehension, interaction and orientation. This is not implying that the causalities need to model exactly the laws of time and motion, but that the "laws" used in the environment need to be continuous throughout that environment, allowing things to be comprehended, and to an extent, explainable. A ball floating in mid air is considered strange, but provide a context of outer space and the ball's behaviour is perfectly acceptable. Attention is given to the issue of causality by Pettifer and West in [Pett97b] and Pettifer in [Pett96].

It is easy to cause navigation and orientation problems if attention is not given to the design of the virtual environment. This would obviously make the system worse than 2D graphics or plain text because the cognitive overload gets so large. Conversely, if suitable attention is paid to the design of the virtual environment, the metaphors used, the interface between the environment and the user, and the use of suitable "laws" (relating to the metaphor if the metaphor allows) then there is a great potential for the use of virtual reality and environments.

3.8 European Virtual Reality Research

The source of this information is now dated but it provides a snapshot of the research in 1994 in European Universities and research centres. Encarnaçao et al. [Enca94], the authors of the article, divided the research into several main categories and discuss an application of each in detail. They also provided a summary of demonstration and (at least at the time) prototype systems. The six main areas identified were

- System Development
- Software Engineering
- Telepresence and Beyond
- Entertainment
- Simulation
- Scientific Visualisation

The application discussed for each of these areas will be briefly summarised, along with the prototype systems mentioned in the paper.

In Germany, IDG (The Fraunhofer Institute for Computer Graphics) developed a VR toolkit known as Virtual Design. This toolkit enables the creation of real time visualisation with integrated audio. Applications such as architecture, interior design, historical reconstruction and CAD have all been demonstrated with this system.

Division Ltd. (based in Bristol, UK) provides both hardware and software for VR systems. They sell a range of accelerator boards for PCs, and for UNIX workstations based on Intel's architecture. They also have dVS, a PC operating environment for VR that sits on top of standard operating systems. They have also ported this software to several UNIX variations.

Advanced Robotics Research Limited (ARRL) is based in Salford, UK. Their research has concentrated on the use of robotics for telepresence. Whilst demonstrations of architecture and scientific visualisation have been produced their research has also involved much to do with robotics such as path planning and obstacle avoidance.

W Industries (which changed to Virtual Entertainment Systems, October 1993) are based in Leicester, UK. They produce VR games, and LegendQuest (an adventure game) deviates from standard gaming ideas because of the use of a "pod". The user has to use a headset and joystick inside the pod and hand and head tracking are used to translate user movement into the gaming environment.

TNO-FEL (The Physics and Electronics Laboratory (FEL) of the Netherlands Organization for Applied Scientific Research (TNO)) research focused on simulation problems. One of their simulations was for astronauts to practise manoeuvring around a space station as they would have to if repairs needed making in space. Again HMDs (Head Mounted Displays) were used with user movement sensing.

GMD (German National Research Center for Mathematics and Computer Science) work on scientific visualisation with VR. They applied visualisation to medicine, fluid flow, molecular modelling and robotics. Research has also concentrated on the enabling graphics technologies to put these visualisations into practice. GMD research into visualisation is continuing and also being published.

There are also several research prototype systems documented including Aviary (A Virtual Reality) created by the Advanced Interface Group at the University of Manchester. Manchester's system is the only one of these that is widely known about in 1998, at least in academic publishing. It has also progressed from Aviary to two systems; MAVERIK [Hubb96], which handles the graphics, and Deva [Pett97a], which deals with the behaviours of objects in the virtual world.

As the authors write in the opening section of the paper, one year is a long time in VR research. Whilst this may have been true in 1994, and to some extents true today, many of these projects and systems are not well known four years on. The games market is certainly not that different. The state of the art (on PCs) can be considered to be games such as Unreal and Quake 2 but generally these benefit only from improved graphics, with better standard hardware. There is not yet the move to the headsets and wands promised by VR hype and marketing. Scientific visualisation has also not progressed that far because of the VR limitations. There is an improvement in the level of graphics available (at least with monitors and projectors) but since VR research is slowing down, the systems and worlds created with it are also slowing in terms of development.

The costs of the specialist pieces of hardware are still too prohibitive and only research centres (both industrial and academic) with large budgets or with generous sponsors can afford to use them. In aiming their systems at people with this technology they are also limiting their market. Manchester have again succeeded in this area because their VR system (using MAVERIK and Deva) can be used with standard monitors, keyboards and mice, but headsets and three dimensional mice can also be utilised.

3.9 Some Existing Scientific Uses

Visualisation is not the only use of VR. In fact to claim this would be extremely short sighted as many other uses of VR have been investigated and proved to be beneficial. Whilst scientific visualisation (of mostly numeric data) has made much use of VR techniques, so have other disciplines such as those involved with training and treatment.

Bryson [Brys96] defines scientific visualisation as

> *"… is the use of computer graphics to create visual images that aid in the understanding of complex (often massive) numerical representations of scientific concepts or results."*

He then lists possible applications of scientific visualisation

- Computational fluid dynamics
- Molecular modelling
- Geological/Astronomical recorded data
- Topological structures

Some of the other applications of visualisation and VR to scientific visualisation are documented in Samtaney et al. [Samt94], Ribarsky et al. [Riba94], Hibbard et al. [Hibb94] and Hanson et al [Hans94]. These papers are in no way a summary of all the scientific visualisation activities being researched, for example VR applied to flow dynamics is a popular research area, but provides a snapshot of areas which tend to be closer to information visualisation.

Hodges et al. [Hodg95] carried out empirical studies using virtual environments for the treatment of the fear of heights. They created an environment that used a real world platform with the environment projected for the user via the use of head mounted displays. They used the platform so that the users had a feeling of realism and were able to grip onto it as they would have done in reality. The *degree of presence* mentioned in *3.5 Immersion* is important for this therapy which is underlined by the authors:

> *"For virtual environments to be effective, they must activate the fear structure and elicit fearful responses. Evoking a sense of presence in a virtual height situation is therefore essential to conducting exposure therapy."*

These studies showed that the use of virtual environments in this situation can be useful. It was also shown that a person's perception and handling of real world situations could be altered by virtual world experiences.

Another medical area that has benefited from the use of VR is the planning of complex neurosurgical operations. The focus of this work by Goble et al [Gobl95] is on preoperative VR rather than as an aid to be used during the operation. The system was developed to allows surgeons to explore the patients head using the information gleaned from magnetic resonance imaging (MRI), digital subtraction angiography images and computed tomography (CT). The system has been tested by neurosurgeons working on real patients and the preliminary results are encouraging.

Yet another use of VR is for simulations, not just on screen using the graphics capabilities of the systems but using physical hardware, and feedback devices to create a responsive environment for the user. Kuhl et al. [Kuhl95] carried out work on a driving simulator. The simulator provides not only visual and motional cues to the driver, but auditory and haptic, in an attempt to create as realistic an environment as possible. The system aims to provide a base for varied scenarios, not just practising or learning to drive. It can be used, for example, to evaluate proposed roads or to test additions to existing roads. It also has a role to play in the testing and evaluation of new features such as automated toll booths or speed limiting devices.

These few examples show the other uses that VR can be applied to, and it may be that the visualisation field can benefit from the advances made in this other research.

3.10 Entertainment

Academia has often been accused of being detached from reality, and such attitudes prompted Potts [Pott93] to write his paper on the practice of using *industry as laboratory* rather than the more popular method of *research then transfer* in the Software Engineering field.

Virtual reality application to the creation of desktop environments can learn much from the gaming industry in terms of design and even human-computer interaction issues. Benford et al. [Benf97b] refer to Doom (created by Id Software, [IdSoftware] now effectively superseded by Quake, and Quake 2) in their paper. They discuss the method of reducing disorientation of users that Doom uses, although for implementation purposes for their work they did not employ such approaches.

> *"We did consider the use of "solid" objects as another way of avoiding disorientation, similar to the approach of games such as Doom where one is constrained by solid boundaries to move through corridors and other enclosing spaces."*

All the games mentioned above (created by Id Software) are of the first person genre where the user's viewpoint of the game is through the eyes of their character and all follow the convention of solid boundaries and what are approximations of real-world physics. Id Software is not the only games company producing such games. A very recent release based on the same principle is Unreal [Unreal], produced by Epic MegaGames. This is a first person shooting game which also supports multiplayer (online) games.

In a panel session held at the 1996 ACM CSCW Conference [Dame96] the following is written

> *"In addition, virtual worlds employ fast 3-D graphic rendering engines found in gaming environments but their application is almost purely social or creative. Avatars do not generally die or kill other avatars in virtual worlds."*

Whilst this killing aspect may be true of many games including Quake 2, there are also many social and teamwork variations of the games. Quake and Quake 2 have a team based variant where the object is to both defend your own flag and try to capture the opponents. The players are split into two teams for this variant and to be successful requires co-operation and communication between players. There is also a great deal of strategy to be employed within the team for aspects such as attack and defence.

Another area where games companies excel is in listening to their customers. They provide alterations and (free) bug fixes more readily than the "serious" software companies, and a lot of these enhancements improve the usability of the system. Over time their products then incorporate these features at release. Many (if not all) games released today for the PC platform provide key, mouse and joystick configuration settings so that the user can create a setup that they are comfortable with, and even have a familiar configurations in each product.

Not only can academia learn from the graphics produced in games, it can also learn from the advances made in human-computer interaction. An added advantage of learning from the concepts used by games companies is that the techniques used have been empirically evaluated by many millions of users.

3.11 Urban Virtual Environments

Urban virtual environments are included on their own because they tie together the concepts of spatial orientation, navigation cues, "reality" and the legibility aspects documented as a way to solve the problems of being lost in hyper/cyberspace.

The five features for legibility documented by Ingram and Benford [Ingr95] are of importance for urban visualisations. The features identified are:

- Landmarks
- Districts
- Paths
- Nodes
- Edges

Landmarks are static and recognisable objects in the environment. Because they are static they can be used as orientation and location markers.

Districts can be defined as areas of the environment that have a local coherence. They can be said to have a distinct character and this allows the entire of a district to be seen as a single item. Several things can identify a section of the environment as a district, common ones being the architecture of the buildings or their use.

Paths can be considered to be main routes of travel through the environment.

Nodes are important items along paths. These can be considered to be like mini-landmarks, in that they provide bearing along the paths.

Edges provide borders to districts or objects, and can be composed of structures or features of the environment.

Ingram and Benford originally applied this work to solving the problem of *getting lost in hyperspace* visualisation of hypertext worlds in their research), but the ideas can be extended into any form of virtual reality.

Dieberger [Dieb93a, Dieb93b, Dieb93c, Dieb94, and Dieb97] carried out navigation and layout work using city metaphors to create The Information City. He used a city metaphor, with houses, to display hypertext information. He ties together metaphors and navigation and writes

> *"When talking about navigation in information spaces we automatically use metaphors but we do not fully use these metaphors. Metaphors are incomplete mappings from a source to a target domain thus carry certain restrictions. We should see these restrictions not as obstacles but as devices to communicate structure to our users."*

and later in the paper

> *"I get the impression that we don't use metaphors to their full potential in navigation. They are not only vehicles to make something easier to understand, but – especially in the case of information spaces – they are also structuring devices."*

If the points made by Dieberger about metaphors are true, then making the metaphor tally with the features identified by Ingram and Benford provides the capability of creating powerful metaphors. These metaphors can then convey not only the information of the information spaces, but provide navigation and orientation information that is based on that information. This means that evolving data sets with suitably evolving visualisations can still remain familiar because the landmarks, and inter-data relationships will evolve with the data thus providing a degree of cognitive familiarity.

4. Conclusions

The areas covered show that there is much more to virtual realities and environments than simply using 3D graphics. Since many of these facets of VR are related to usability then it is important that they are considered properly by people developing VR environments.

VR has the potential to be very useful to the visualisation field as it provides an interactive 3D environment in which to build the visualisations. This allows the visualisation field to move into new areas, previously inaccessible when using only 2D graphics. Another advantage is the increasing availability of VR systems, especially for platforms such as the PC, hardware it is reasonable to expect on a user’s desk in industry.

Overall VR has the potential to be a very useful tool, although just relying on its usage is not enough. Careful thought and good design need to be employed to allow it to reach it’s full potential.

8. References

[Auks92] S. Aukstakalnis, D. Blatner.
Silicon Mirage: The Art and Science of Virtual Reality.
Peach Pit Press, 1992. ISBN 0-938-151-82-7.

[Benf96] S. Benford, C. Brown, G. Reynard, C. Greenhalgh.
Shared Spaces: Transportation, Artificiality and Spatiality.
Proceedings of the 1996 ACM Conference on CSCW, Boston, Masachusetts, USA, November 16-20, 1996, pp77-85. ISBN 0-89791-765-0.

[Benf97b] S. Benford, C. Greenhalgh, D. Snowdon, A. Bullock.
Staging a Public Poetry Performance in a Collaborative Virtual Environment.
Proceedings of the 5th European Conference on CSCW '97, pp125-140, Kluwer Academic Publishers, ISBN 0-7923-4638-6.

[Bent93] R. Bentley, T. Rodden, P. Sawyer, I. Sommerville.
Architectural Support for Co-operative Multi-User Interfaces.
Research Report CSCW/11/93, Centre for Research in CSCW, Lancaster University, 1993.

[Bola94] M. T. Bolas.
Human Factors in the Design of an Immersive Display.
IEEE Computer Graphics and Applications, pp55-59, January 1994.

[Bowe96] J. Bowers, J. O'Brien, J. Pycock.
Practically Accomplishing Immersion: Cooperation in and for Virtual Environments.
Proceedings of the 1996 ACM Conference on CSCW, Boston, Masachusetts, USA, November 16-20, 1996, pp380-389. ISBN 0-89791-765-0.

[Brys96] S. Bryson.
Virtual Reality in Scientific Visualization.
Communications of the ACM, Vol. 39, No. 5, pp62-71, May 1996.

[Card91] S. K. Card, G. G. Robertson, J. D. Mackinlay.
The Information Visualizer: An Information Workspace.
Xerox Palo Alto Research Centre, Palo Alto, California 94304, 1991.

[Cast97] J. Castellanos.
Postulates and Proofs.
On-line document available from http://riceinfo.rice.edu/projects/NonEuclid/proof.html.

[Chur98] E. F. Churchill, D. Snowdon.
Collaborative Virtual Environments: An Introductory Review of Issues & Systems.
Virtual Reality, Vol. 3, No. 1, pp3-15, 1998.

[Dame96] B. Damer, A. Bruckman (Moderators).
PANEL: Peopled Online Virtual Worlds: A New Home for Cooperating Communities, a New Frontier for Interaction Design.
Proceedings of the 1996 ACM Conference on CSCW, Boston, Masachusetts, USA, November 16-20, 1996, pp441-442. ISBN 0-89791-765-0.

[Davi95] S. Davis.
A Guessing Measure of Program Comprehension.
International Journal of Human-Computer Studies, Vol. 42, pp245-263, 1995.

[Dieb93a] A. Dieberger.
The Information City - A Step Towards Merging of Hypertext and Virtual Reality.
Poster at Hypertext 1993.

[Dieb93b] A. Dieberger.
The Information City - A Metaphor for Navigating Hypertexts.
Presented at BCS-HCI '93, Loughborough, England, September 1993.

[Dieb93c] A. Dieberger, J. G. Tromp.
The Information City Project - A Virtual Reality User Interface for Navigation in Information Spaces.
Was to appear in Proceedings of the Symposium Virtual Reality Vienna, Vienna, December 1-3, 1993 but was never printed.

[Dieb94] A. Dieberger.
Navigation in Textual Virtual Environments using a City Metaphor.
PhD thesis, Vienna University of Technology, Faculty of Technology and Sciences, November 1994.

[Dieb97] A. Dieberger.
Navigation Metaphors and Social Navigation in Information Spaces.
Position Paper for the CHI '97 workshop on Navigation in Information Spaces, 1997.

[Elli94] S. R. Ellis.
What Are Virtual Environments?
IEEE Computer Graphics and Applications, pp17-22, January 1994.

[Enca94] J. Encarnaçao, M. Göbel, L. Rosenblum.
European Activities in Virtual Reality.
IEEE Computer Graphics and Applications, pp66-74, January 1994.

[Find95] J. M. Findlay, F. N. Newell.
Perceptual cues and object recognition.
In "Simulated and Virtual Realities: Elements of perception", edited by K. Carr and R. England, Taylor & Francis Publishers, 1994, pp113-130. ISBN 0748401288.

[Fitz96] G. Fitzpatrick, S. Kaplan, T. Mansfield.
Physical Spaces, Virtual Places and Social Worlds: A study of work in the virtual.
Proceedings of the 1996 ACM Conference on CSCW, Boston, Masachusetts, USA, November 16-20, 1996, pp 334-343. ISBN 0-89791-765-0.

[Gard93] B. R. Gardner.
The Creator's Toolbox.
In "Virtual Reality: Applications and Explorations", A. Wexelblat (Ed.), Academic Press Professional, Cambridge, MA, 1993, pp91-121. ISBN, 0-12-745045-9.

[Gibs84] W. Gibson.
Neuromancer
Paperback edition publiched in 1995 by Voyager, an imprint of HarperCollins Publishers. ISBN 0 00 648041 1

[Gobl95] J. C. Goble, K. Hinckley, R. Pausch, J. W. Snell, N. F. Kassell.
Two-Handed Spatial Interface Tools for Neurosurgical Planning.
Computer, pp20-26, July 1995.

[Hans94] A. J. Hanson, T. Munzner, G. Francis.
Interactive Methods for Visualizable Geometry.
Computer, pp73-83, July 1994.

[Hemm94] M. Hemmje, C. Kunkel, A. Willett.
LyberWorld - A Visualization User Interface Supporting Fulltext Retrieval.
Proc. of SIGIR '94, Dublin, 1994.

[Hibb94] W. L. Hibbard, B. E. Paul, D. A. Santek, C. R. Dyer, M-F. Voidrot-Martinez.
Interactive Visualization of Earth and Space Science Computations.
Computer, pp65-72, July 1994.

[Hodg95] L. F. Hodges, R. Kooper, T. C. Meyer, B. O. Rothbaum, D. Opdyke, J. J. de Graaff, J. S. Williford, M. M. North.
Virtual Environments for Treating the Fear of Heights.
Computer, pp27-33, July 1995.

[Hubb93] R. Hubbold, A. Murta, A. West, T. Howard.
Design Issues for Virtual Reality Systems.
Presented at the First Eurographics Workshop on Virtual Environments, Barcelona, 7th September 1993.

[Hubb96] R. J. Hubbold, X. Dongbo, S. Gibson.
MAVERIK - the Manchester Virtual Environment Interface Kernel.
Proceedings of the 3rd Eurographics Workshop on Virtual Environments, Monte-Carlo, February 1996.

[IdSoftware] Id Software homepage can be found online at http://www.idsoftware.com/.

[Ingr95] R. Ingram, S. Benford.
Legibility Enhancement for Information Visualisation.
Proceedings of Visualization 1995, Atlanta, Georgia, October 30th - November 3rd 1995.

[Isda93] J. Isdale.
What is Virtual Reality?
On-line document available from http://www.cms.dmu.ac.uk/~cph/VR/whatisvr.html.

[Kuhl95] J. Kuhl, D. Evans, Y. Papelis, R. Romano, G. Watson.
The Iowa Driving Simulator: An Immersive Research Environment.
Computer, pp35-41, July 1995.

[Levi95] S. Levialdi, A. Massari, L. Saladini.
Visual Metaphors for Database Exploration.
A position paper relating the Virgilio system, available on-line from http://www.darmstadt.gmd.de/~hemmje/Activities/Virgilio/Publications/virgilio.ps.

[Mach94] C. Machover, S. E. Tice.
Virtual Reality.
IEEE Computer Graphics and Applications, pp15-16, January 1994.

[Pesc93] M. D. Pesce.
Final Amputation: Pathogenic Ontology in Cyberspace.
Presented at the Third International Conference on Cyberspace, University of Texas, May 1993.

[Pett96] S. R. Pettifer.
Beyond the Metaphysics of Virtual Reality.
M. Sc. Thesis, Department of Computer Science, University of Manchester, UK. January 1996.

[Pett97a] S. Pettifer, A. West.
Deva: A coherent operating environment for large scale VR applications.
Presented at the first Virtual Reality Universe conference in Santa Clara, California, April 1997.

[Pett97b] S. Pettifer, A. West.
Worlds within Worlds: on the topology of virtual space.
Presented at Cyberconf, Oslo (http://www.cyberconf.org), June 1997.

[Pott93] C. Potts.
Software Engineering Research Revisited.
IEEE Software, pp19-28, September 1993

[Riba94] W. Ribarsky, E. Ayers, J. Eble, S. Mukherjea.
Glyphmaker: Creating Customized Visualizations of Complex Data.
Computer, pp57-64, July 1994.

[Samt94] R. Samtaney, D. Silver, N. Zabusky, J. Cao.
Visualizing Features and Tracking Their Evolution.
Computer, pp20-27, July 1994.

[Sand97] O. Sandor, C. Bogdan, J. Bowers.
Aether: An Awareness Engine for CSCW.
Proceedings of the 5th European Conference on CSCW '97, pp221-236, Kluwer Academic Publishers, ISBN 0-7923-4638-6.

[Step92] N. Stephenson.
Snow Crash
Paperback edition published by ROC (Penguin Group) 1993. ISBN 0-14-023292-3.

[Unreal] The official Unreal web site can be found at http://www.unreal.com/.

[Ware95] C. Ware.
Dynamic Stereo Displays.
Proceedings of SIGCHI '95, ACM, Denver, Colorado, May 7-11, 1995. Available at http://www.acm.org/sigchi/chi95/proceedings/papers/cw_bdy.htm.