"The World Is What You Make It" - An Application Of Virtual Reality To The Tourism Industry.

Submitted by

Patrick P. Horan
B.Sc.(Mgmt).
H.DIP (Hotel and Catering Management),
MIHCI

in fulfillment of the requirements for a Masters Degree

To

Dublin City University,
Dublin 9.

Supervisor of Studies

Ciaran MacDonaill, Ph.D.
Dublin Institute of Technology,
Cathal Brugha St.

April 1997.

Declaration

I hereby certify that this material, which I now submit for assessment on the programme of study leading to the award of a Masters Degree, is entirely my own work and has not been taken from the work of others save and to the extent that such work has been cited and acknowledged within the text of my own work.

Patrick Horan (DIT 039/97)

Signed Patrick Horan
Candidate

Date 25th June 98

Abstract

The tourism industry is a highly information intensive-industry. In few other areas of activity are the generation, gathering, processing, application and communication of information as important for day-to-day operations as they are for the tourism industry (Buhalis 1994). Traditional sources of tourism information, images, text, sound, animation and video, provide potential tourists with short and often rather limited glimpses of tourism destinations which may be inadequate to enable them to make informed decisions (Cheong 1995). In addition, these sources of tourist information provide only a passive experience as they often possess little involvement on the part of the potential tourist. Virtual Reality (VR), on the other hand, enables potential tourists to interact with and experience each tourist destination in high detail and provides them with enough information to make a well-informed tourist decision. When considering its application within the tourism industry, VR will offer the ability not only to view a destination, but also, to participate in the activities offered at the destination. Through VR the tourist advances from being a passive observer to being an active participant (Williams & Hobson 1994).

This thesis addresses issues associated with the design and evaluation of a VR application to the tourism industry that provides users with all the traditional types of tourist information along with allowing them to experience a multi-participant, realistic, interactive and real-time walkthrough of real-life tourist destinations. In order to develop these walkthroughs, the basic concepts of VR had first to be analysed. This was achieved by gaining hands-on experience of the different types of VR hardware and software available in conjunction with an in-depth literature review. Following the completion of this analysis, an overview of the tourism industry was developed. This overview identified certain properties of the tourism product that lend themselves readily to the application of VR. Once this was completed the final stage of the research was concerned with the development of the walkthroughs and the elicitation of knowledge from the development of these walkthroughs.

There were many conclusions uncovered by this research but the most important was that VR can indeed be applied successfully to the tourism industry. The main areas of application will be in the areas of tourism policy and planning and the marketing of the tourism product. Another conclusion that was drawn from this research was that VR applications can help to generate realistic impressions and expectations of what can be experienced at a tourism location. The final outstanding conclusion drawn from this research was that potential tourists viewed the VR application as a decision making tool that increases their desire to actually visit a tourist location and not as a tourism substitute.

Table of Contents

Chapter Five - The Design Process - The Making of D-VR3.

Chapter Six - Evaluation of The Virtual Environments.

Chapter Seven - Recommendations and Conclusions.

Acknowledgments

The author would like to acknowledge the following for their assistance in the completion of this research.

This dissertation would simply never have happened without two years of constant encouragement and continual support from, his supervisor, Dr. Ciaran MacDonaill. Ciaran has remained excited about the ideas that have driven our research and about the sometimes strange results that we have discovered. As both a friend and a mentor, he has been invaluable to me over the years. I offer my thanks for showing me the fun approach to academia, and for getting me involved in the publish-or-perish process from the beginning. I look forward to working with him in the future.

To Peter O'Connor for his continuous help and guidance throughout the course of this dissertation.

I would like to extend my deepest gratitude to my parents, for their love, patience, dedication, and support all my life.

To Siemens Nixdorf for their continuous support throughout the course of this entire dissertation.

To Neil, Tom and Gerry for putting up with the sometimes off-the-wall requests and questions.

The librarians of DIT Cathal Brugha St., for their constant help and seemingly endless patience.

To Roseanne, Joe, Johnny, Colin, Niall and Hubert for..................................I'll think of something.

Thanks to Neil for the inspiration.

To my Dublin family: a big thanks to Frank for the alarm calls, for always being bright and cheery in the mornings and, of course, for being both my dictionary and thesaurus. Seriously, Thanks Lang. And most of all, to Glenda for showing me the light and for being by my side every step of the way....

Patrick P. Horan
April 1997

List of Illustrations

List of Tables

List of Abbreviations

2D	Two Dimensional.
3D	Three Dimensional.
6D	Six Dimensional.
3DOF	Three Degrees of Freedom.
5DOF	Four Degrees of Freedom.
6DOF	Six Degrees of Freedom.
ACS	Action Code Script.
ARPA	Advance Research Projects Agency.
AVRIL	A Virtual Reality Interface Library.
BOOM2	Binocular Omni-Orientation Monitor 2.
BRENDER	Blazing Render.
CAD	Computer Aided Design.
CAT	Computerised Axial Tomography.
CD-ROM	Compact Disc - Read Only Memory.
CRS	Central Reservations Systems.
CRT	Cathode Ray Tube.
CSG	Computer Solid Geometry.
DEU	Doom Editor Utility.
DIT	Dublin Institute of Technology.
DVP	Digital Video Producer.
D-VR3	Desktop Virtual Reality Version 3.
DXF	Data Exchange Format
FOV	Field of View / Vision.
FPS	Frames per second.
GNP	Gross National Product.
GUI	Graphical User Interfaces.
HETH	Hexen Editor For Total Headcases.
HMD	Head Mounted Display.
I/O	Input/Output.
IT	Information Technology.
LAN	Local Area Networks.
LCD	Liquid Crystal Display.
LED	Light Emitting Diode
MIT	Massachusetts Institute of Technology.
MRI	Magnetic Resonance Imaging.
OPW	Office of Public Works.
PC	Personal Computer.
POLY/SEC	Polygons per second.

RSI	Repetitive Strain Injury.
SGI	Silicon Graphics Interface.
SPSS	Statistical Package for Social Sciences.
SRS	Sound Retrieval System.
TIS	Tourist Information System.
UNC	University of North Carolina.
VE	Virtual Environment.
VPL	Visual Programming Language.
VR	Virtual Reality.
VRML	Virtual Reality Modelling Language.
VROOM	Virtual Reality Using Object-Oriented Methods.
WAN	Wide Area Networks.
WTK	World Toolkit.
WTO	World Tourism Organisation.
WTTC	World Travel & Tourism Council.

Trademarks

3D Studio is a registered trademark of Autodesk Software.

Animation Player is a registered trademark of Autodesk Software.

Corel Photo-Paint Version 6.0. is a registered trademark of Corel Software.

Crystal Eyes is a registered Trademark of Stereographics Incorporated.

DataGlove is a registered trademark of Visual Programming Languages.

Digital Video producer is a registered trademark of Asymetrix Corporation.

Doom is a registered trademark of Id Software.

Duke Nukem is a registered trademark of 3D Realms Software.

Editart is a registered trademark of 3D Realms Software.

Heretic is a registered trademark of Raven Software.

Hexen is a registered trademark of Raven Software.

MRG2 is a registered trademark of Liquid Image.

MS-DOS is a registered trademark of the Microsoft Corporation

Paperport is a registered trademark of the Visioneer Corporation.

PowerGlove is a registered trademark of the Mattel Corporation.

QuickTime VR is a registered trademark of the Apple Corporation.

Sound Blaster AWE32 is a registered trademark of Creative Laboratories.

Sound Retrieval System is a registered trademark of SRS Laboratories.

SPSS is a registered trademark of SPSS Incorporated.

Toolbook Version 3.0. is a registered trademark of Asymetrix Corporation.

Virtus Walkthrough is a registered trademark of the Virtus Corporation.

Windows is a registered trademark of the Microsoft Corporation.

Wolfenstein is a registered trademark of Id Software.

World Toolkit is a registered trademark of Sense8 Corporation.

List of Appendices

Introduction

Research Outline

The purpose of this research is:

1. To explore the use of Virtual Reality (VR) in order to develop *multi-participant*, *realistic*, *interactive* and *real-time* walkthroughs on commercially available Personal Computers (PCs). These walkthroughs could be applicable to any industry; however, this research concentrates on the tourism industry.

2. To incorporate these walkthroughs into a Tourism Information System (TIS) with the use of a multimedia-based environment. This TIS will provide potential tourists with additional information to make conscious, well-informed tourism decisions.

Rationale

Traditionally, the tourism industry is a highly information intensive-industry. In few other areas of activity are the generation, gathering, processing, application and communication of information as important for day-to-day operations as they are for the tourism industry (Buhalis 1994). Unlike durable and industrial goods, the intangible tourism product cannot be displayed or inspected at the point of sale before it is purchased. Furthermore, the tourism product is normally bought long before the time of use and away from the place of consumption. Information is the way the product is presented to the potential tourist. Consequently, tourists require a wide variety of specific regional information such as accessibility, facilities, attractions, and local activities. The provision of timely and accurate information relevant to the potential tourist's needs is often the key to successful satisfaction of the tourism demand (Williams 1993).

Largely due to the problems of dealing with an increasing volume of potentially relevant information, the tourism industry is becoming increasingly reliant upon, and influenced by, developments in Information Technology (IT) (Bennett & Radburn 1991). The unique features of VR provide a novel way of presenting the travel and tourism product (Cheong 1995) and it is, therefore, conceivable that VR could have a substantial effect on the tourism industry in the future. From a marketing perspective, VR has the potential to revolutionise the promotion and selling of the tourism product (Hobson & Williams 1995). Tourism information is currently being dispersed by the use of travel agents, videos, travel brochures and tourism multimedia software. These traditional types of tourism information media provide prospective tourists with only short and rather limited glimpses of a destination's attractions (Cheong 1995). At times, the accompanying destination information may be inadequate for clients to make travel decisions. VR, in comparison, provides clients with the opportunity to explore each destination in great depth prior to making their travel arrangements. The client would, therefore, be provided with ample information to allow them to form realistic expectations of the travel destination prior to their visit. In all, VR could be used to dispel elements of uncertainty and ensure that the visitor's expectations of his impending visit are equitable to the subsequent actual experience (Cheong 1995).

Objectives

The objectives of this research are as follows:

1. Conduct an investigation in order to determine the VR applications currently available.

2. Conduct a comprehensive analysis of the VR applications available on the PC.

3. Compare and contrast the VR applications available on the PC.

4. Develop prototype walkthroughs of an area of tourism interest using appropriate applications.

5. Compare and contrast the walkthroughs developed using a strict set of evaluation criteria.

6. Incorporate the Walkthroughs developed into a Multi-Media environment.

Methodology

In order to develop *multi-participant*, *realistic*, *interactive* and *real-time* walkthroughs of areas of tourism interest, the basic concepts of Virtual Reality have first to be analysed. This is to be achieved by using a two-tier approach. The first tier consists of an in-depth literature review of available textbooks, journals, conference papers and web sites. The second tier involves gaining experience of using the different types of VR hardware and software currently available. The objective of this two-tier approach is to gain a firm understanding of the requirements associated with the development of VEs.

Following the specification of these requirements, an overview of the tourism industry is developed from a literature review of the tourism product and the structure of the tourism industry. This overview identifies certain properties of the tourism product that lend themselves readily to the application of VR.

This dissertation is firstly concerned with the development of virtual walkthroughs of areas of tourism interest and their incorporation into a TIS and, secondly, with the elicitation of knowledge required from the development of such a system. This elicitation is conducted in three ways:

1. The evaluation of each of the VEs created against a strict set of performance criteria based on the environment's *interactivity, vividness* and *time*, Section 6.2. This process is used to identify how effectively each environment copes with each of the criteria used and its effectiveness as a source of tourism information.

2. The second form of knowledge elicitation concerns an investigation using peer reviews. The peer reviews are composed of individuals from the fields of computers, architecture and from the tourism industry. These discussions are invaluable in that they encourage the group

members to discuss topics concerned with the specific system developed and this facilitates the development of a more effective and user-friendly source of tourism information.

3. The third, and final, form of knowledge elicitation involves the use of a comprehensive questionnaire process. This process involves administering a survey to potential tourists in an attempt to determine how such an TIS compares to traditional sources of tourism information and also to facilitate the development of a more effective system.

The knowledge gained from the above methods serves to evaluate each Virtual Environment (VE) from a potential tourists' perspective in order to determine its effectiveness as a source of tourist information and to help create a process for the evaluation of past, present and future VEs.

Chapter Outline

Chapter One introduces the basic concepts involved in VR and highlights areas such as choosing an appropriate definition for this particular research, the evolution of VR, the construction of an appropriate model and the dimensions of VR. The final part of this chapter examines the concepts of multimedia and the synergy that exists between VR and multimedia.

Chapter Two concentrates on the specific hardware and software components which form the basis for VR systems at present. The three broad types of VR system presently available are then discussed in an attempt to chose the most appropriate to the tourism industry.

Chapter Three examines the different commercial applications of VR. These primary applications are in entertainment, architecture, medicine, education/training and manufacturing. Specific applications vary from flight simulators for NASA's space shuttle program to arcade games.

Chapter Four addresses the issues associated with the introduction of VR to the tourism industry. The first part of this Chapter discusses the tourism industry, its relationship with IT, and the possible areas where VR applications are most likely to occur in the tourism industry in the future. The second part of the Chapter addresses the considerations associated with the design of D-VR3, the specific application to the tourism industry constructed in this dissertation.

Chapter Five describes the design process used in the construction of the different VEs used in D-VR3. This chapter is divided into two sections. The first section describes the steps involved in the construction of the Prototype model (the Computer Section, DIT, Cathal Brugha St.), and the second section describes the construction of the Actual model (the Monastic Enclosure, Glendalough). Both the Prototype and the Actual models were designed using a variety of software packages varying from three-dimensional (3D) graphics and animation packages to the engines used in 3D games. Finally the process used to incorporated each of these environments into a multimedia-based environment is described.

Chapter Six describes the evaluation of each of these VEs using an internal and external evaluation process. The internal evaluation is a process which involves the evaluation of the specific features of the VE against a strict set of criteria. The external evaluation process involves the evaluation of the VEs using methods such as a questionnaire process and peer reviews.

Chapter Seven concentrates on the conclusions drawn from the findings of both the primary and secondary research undertaken and recommendations are made concerning the role of VR in the tourism industry.

Chapter One

Defining Virtual Reality

1.1. Introduction.

The concept of Virtual Reality (VR) is perceived by many to be in its infancy but the origins of VR can be traced as far back as 'The Ultimate Display', a paper written by Ivan Sutherland in 1965. In this paper Sutherland issued the following challenge: "The screen is a window through which one sees a Virtual World. The challenge is to make that world look real, act real, sound real and feel real" (Sutherland 1965). Over the past decade this challenge has become the research agenda for a rapidly growing community of researchers and industries. Although VR technology has been slowly developing over this period the possibilities inherent in the new medium have only recently crossed a cultural threshold (Biocca 1992): that VR has begun to shift away from the purely theoretical and towards the practical. It is the implications of this orientation towards the practical applications of VR that are important for the tourism industry.

1.2. Definition - What is Virtual Reality

As a new technology seeks its identity, there is bound to be ambiguity in early attempts to define it. Nonetheless, in an attempt to provide a complete understanding of VR, its applications and concepts, one must first give a comprehensive definition of VR. There is no standard definition of VR. However, in an attempt to formulate a workable definition the three main approaches to the VR definition must be introduced. These three approaches may be categorised by the type of technology used, the type of personal response that the participant experiences and, the type of Virtual Environment (VE) that is created.

1.2.1. Definition in Terms of Technology.

The most common type of definition used is based upon technology and it pays particular attention to the component parts of an "Immersive" technological VR system. This system usually includes a computer capable of real-time animation, controlled by a set of wired gloves and a position tracker, and using a head-mounted stereoscopic display for visual output. The following definition is an example of a technical definition:

> *Virtual Reality is electronic simulations of environments experienced via head-mounted eye goggles and wired clothing enabling the end-user to interact in realistic three dimensional situations (Coates 1992).*

This type of definition includes the concept of artificially created environments and immersive input devices as the means to access these environments. The application of this type of definition is limited to the technologies that it describes; its unit of analysis and its potential for variance are left unspecified (Steuer 1992). Therefore, if one were to employ this type of definition, one would be clearly limiting the scope of VR and its applications.

1.2.2. Definition in Terms of Presence.

This type of definition describes the participant's personal response to VR. Typically such an approach is based on the experiences or the feelings of the participant. As the description of VR focuses more on experience and less on interaction, the concept of passive participation must be addressed. The following definition is an example of a definition in terms of presence:

> *Virtual Reality is a remote and artificially constructed environment in which one feels a sense of presence, as a result of using a communication medium (Steuer 1992).*

The key to defining VR in terms of human experience rather than technological hardware is a concept called *presence*. Presence can be thought of as the experience of one's physical environment; it refers not to one's surroundings as they exist in the physical world, but to the perception of those surroundings as mediated by both automatic and controlled mental processes (Gibson 1979). Presence is defined as the sense of being in an artificial or remote environment (Sheridan 1992). Many perceptual factors help to generate this sense, including input from some or all sensory channels, as well as more mindful attentional, perceptual, and other mental processes that assimilate incoming sensory data with current concerns and past experiences (Gibson 1966). The advantage of defining VR in terms of presence is that by employing the concept of presence, VR can now be defined without reference to particular hardware systems but problems arise because this can be a difficult type of definition to handle as it is relative to each individual (Latta 1991).

1.2.3. Definition in Terms of Virtual Environments.

This type of definition is based on the creation of the VE. There are many ways that this can be accomplished. This definition covers all aspects of VR without discussing the participants' response or the technologies used.

> *Virtual Reality is a human-computer interface where the computer and its devices create a sensory environment that is dynamically controlled by the action of the individual so that the Virtual Environment appears real to the participant (Latta 1991).*

The devices discussed in previous definitions (computers, position sensors, head-mounted displays, etc.) are all relatively recent developments and, with the rate of current developments in technology may, in the future, become obsolete. However, the definition of VR in terms of Virtual Environments (VEs) can be applied to past, present, and future media technologies. Thus, this definition of VR provides a framework in which such newly developed technologies can be examined in relation to past, present and future technologies.

Furthermore, defining VR in terms of VEs alleviates three difficulties that the other two approaches to VR definition fail to cope with successfully. Firstly, this definition of VR refers to the VE, rather than to specific devices and systems used to create that environment. This type of definition shifts the focus of VR away from the type of hardware that is used to create the VE and onto the VE that is actually created. Secondly, this definition specifies the unit of analysis of VR as being the VE rather than the individual, as it is when VR is defined in terms of presence. Therefore, the definition is less subjective, since it does not consist of an individual's experience of presence. Finally, since this definition is not based on technology, it permits variation across technologies along a number of dimensions.

This thesis is not primarily concerned with the technology available to produce "all-inclusive, immersive" VR models, nor is it primarily concerned with generating some sense of "presence". The focus is on building realistic, real-time, interactive models of areas of tourism interest. Consequently the definition of VR in terms of the VE is adopted for the purposes of this research.

1.3. The Evolution of Virtual Reality - A Slow Progression or a Paradigm Shift?

New media are often a small refinement on existing technologies or paradigmatic shift in the way that information is presented and absorbed. VR is an integrative technology - it uses several unrelated technologies to achieve more than the simple sum of its components. VR is a "technology cluster" (Rogers 1993). What this means is that VR is not composed of just one technology, but instead it is composed of a synergistic combination of a host of technologies (Biocca 1992). Computer graphics, position tracking, real-time interaction, and stereo displays have all existed for many years, but as isolated technologies. When these technologies are coupled together they create the illusion of a self-contained, apparently real environment. In order to trace the origin and development of VR one must first of all trace the evolution of its component technologies.

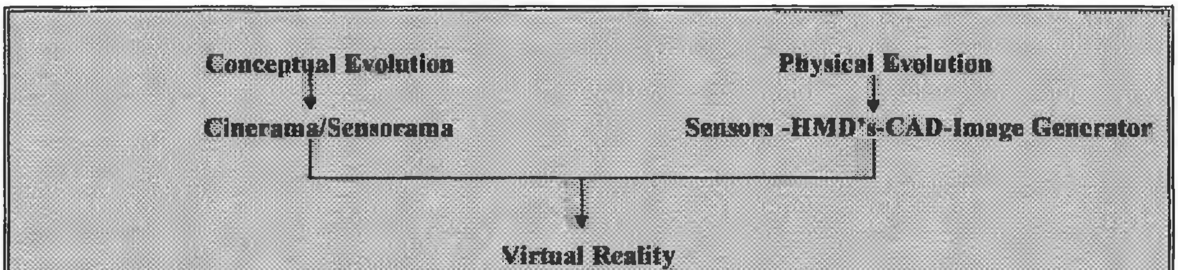

Figure 1.1. The Two Tier Approach to Virtual Reality Development.

The evolution of VR can be best described by adopting a two-tier approach, Figure 1.1. The first tier describes the conceptual evolution of VR from its inclusion in early film-making techniques up to its present computer generated all-inclusive VEs. The second tier discusses the physical evolution of its component parts from their individual invention to their present day convergence.

From a conceptual point of view the development of VR was influenced by film techniques such as stereoscopic, or three-dimensional (3D), cinema and several wide-screen systems that Hollywood filmmakers were experimenting with during the early 1950's. Cinerama, the best-known of these technologies, sought to expand the movie-going experience by filling a larger portion of the audience's visual field. This effect was achieved by the use of three cameras, shooting from slightly different angles. The film was then synchronised and projected onto three large screens that curved inward, wrapping around the audience's peripheral visual field. While the technology proved too costly to be adopted by most commercial theatres, the concept of visual immersion would go on to become an important VR element.

Cinerama excited a young documentary filmmaker named Morton Heilig, who believed the future of cinema lay in creating films that could employ the human senses of sight, sound, scent, and feel. He fused the elements he felt were necessary to create that total illusion, such as the brain's sensory channels and the body's motor network. He called his end product "experience theatre." Heilig's research led to "Sensorama," a VR-type arcade attraction he designed and patented in 1962. Sensorama simulated all the sensory experiences of a motorcycle ride by combining 3D movies, stereo sound, wind, and aromas. By gripping the handlebars on a specially equipped motorcycle seat and wearing a binocular-like viewer, the "passenger" could travel through scenes including California sand dunes and Brooklyn streets. Small grilles near the viewer's nose and ears emitted breezes and authentic aromas. Although Sensorama was extremely complex for the arcade environment, and funding never materialised for a simplified version, Heilig's vision of a medium that combined multi-sensory artificial experiences is fast becoming a reality.

The roots of VR, at least in its popular conception, can be traced back to a thesis presented by a graduate student named Ivan Sutherland in the early 1960s. Sutherland believed that display screens and digital computers could offer a means of gaining familiarity with concepts not realisable in the physical world by providing a window, or looking glass of sorts, into the mathematical wonderland of a computer. Sketchpad, the program Sutherland developed and described in his thesis, used computer technology to create images from abstract ideas. Using Sketchpad and a pen-like device, a computer could create sophisticated images on a display screen resembling a television set. The system responded by rapidly updating the drawing so that the relationship between the user's action and the graphical display was clear. Computer-aided design (CAD) grew out of Sutherland's thesis and has become one of the most powerful components of VR development in the 1990's.

Sutherland next focused on developing technology that would allow computer users to actually enter the world of computer-generated graphics. In 1965, with support from the Department of Defence's Advance Research Projects Agency (ARPA) and the Office of Naval Research, Sutherland unveiled the first system developed to surround people in 3D displays of information, the head-mounted display (HMD). The HMD has become a cornerstone of Immersion VR technology, discussed in more detail in section 2.4.3.

In the late 1960's and 1970's, research on a number of fronts formed the basis of VR as it appears today. Projects such as the Aspen Movie Map, developed by a group of researchers including Andrew Lippman, Michael Naimark and Scott Fisher at the Massachusetts Institute of Technology (MIT), showed video images of Aspen, Colorado, that visitors could actually navigate by indicating their choices on a touch-sensitive display screen. Videoplace, one of several experimental artistic environments designed by arts scholar Myron Krueger, used computers to create what Krueger called "artificial reality," allowing viewers to interact with computer-generated graphics and projected images.

For many of the years since these early experiments VR was all but lost with few interested parties other than the military. In the military, much money has been spent on developing VR displays as flight simulators, and for decades this was the most common commercial application of VR. That situation changed in the mid-eighties when, due to the development of high performance computers the different technologies that enabled the development of VR converged to create the first true VR system.

In 1987, Jaron Lanier, director and founder of Visual Programming Language (VPL), coined the phrase "Virtual Reality" to bring all of the virtual projects under a single term. Over the past decade there have been huge advances in technology including great improvements in three areas that are particularly critical to independent VR research (Larijani 1994):

1. Liquid Crystal Displays (LCD) and Cathode Ray Tube (CRT) display devices;
2. Image generating Systems (computer systems developed to produce appropriate sensory stimuli); and
3. Tracking systems (for converting position and orientation information into computer readable signals that can be reflected in images).

It was then only a matter of time before VR programs began to appear, ranging from VR theme parks to operating rooms, largely aided by products developed by Jaron Lanier, whose programming language operated the first data glove at the NASA research centre. Lanier's company was the first company to focus its efforts on developing products for the infant VR industry, and provided the headsets and gloves used in many early VR applications.

As previously mentioned, VR is not a radically new idea; rather it differs only in degrees from previous systems. It all sounds rather straightforward, but creating the artificial 3D environment that makes up the "other world" has only recently been possible. Wrapping pictures and sound around us and immersing our senses in such a way that the line between the real and illusionary worlds disappears is being done only through a dynamic convergence of many different technologies, each of which evolves and matures at its own pace. Computer speed and power have had to be combined with advances in image processing, tracking mechanisms and intuitive human-to-computer communication to converge into the experiential medium called VR. Previously, conventional computers simply did not work fast enough to display high quality images quickly. There was always a trade-off between time or essence. A computer designed to display high quality images in quick succession has to have both computing power and high speed, as well as good display features. At any given time these technologies are at different stages of development. Interplay among them always involves costly trade-offs in time or quality. In order to create Virtual Worlds, each one of the technologies involved must have reached a stage at which its strength and resources can be used effectively with the other technologies. Another factor that must be taken into consideration when analysing VR is that its diffusion into society will also depend on the diffusion rates of its component parts.

VR technology still has the look and feel of a prototype, just a portal looking out on a more mature technology to come. When the technology is fully implemented and diffused, the medium could be the catalyst for a revolutionary change in the way we communicate.

1.4. Developing a Virtual Reality Model.

The general definition of VR in terms of the VE can be reduced to a general systems model. This model is clearly illustrated in Figure 1.2. The model is divided into two parts - Actual Reality, on the left, and Virtual Reality on the right. The model is described from a participant's perspective. The individual's actions are recorded by the VR system by way of the Interface and the environment is manipulated accordingly.

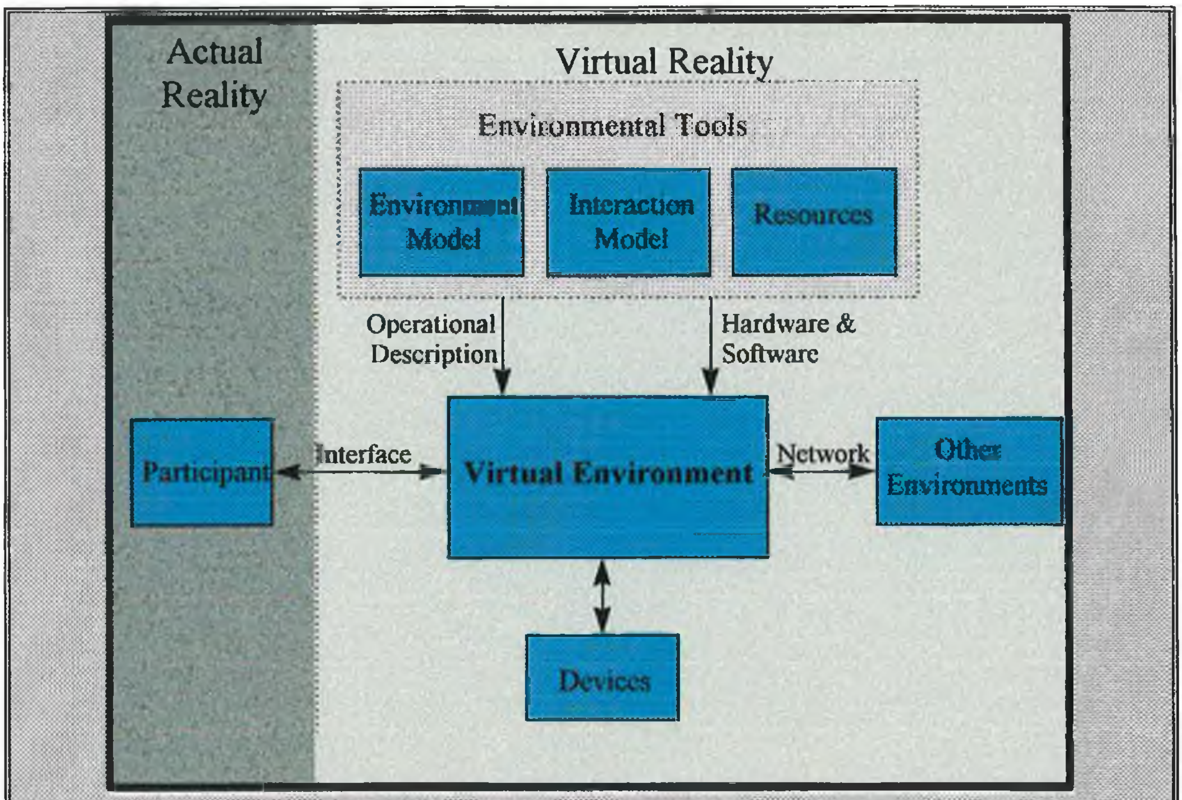

Figure 1.2. The Virtual Reality Model.

The VE is created using a set of tools, shown as the Environment tool at the top of the model and include:

1. An Environment model, which outlines the parameters of the basic VE;
2. An Interaction model, which governs what objects interact and how they interact with the individual and with other objects;
3. A set of resources to support these tools.

The VR system then forms and modifies the VE that is supplied to it. The tools can also take the form of hardware and software that supply any of the components which describe the environment. This model is intended to describe the basic framework in which the system can be implemented and, thus, not every system will contain every element of this model.

To illustrate, the environment model could be the design of a tourist destination using some form of 3D computer-aided visualisation package. The interaction model could be what would happen if the user attempts to open a door or interact with an object in the model. The set of resources to support these tools would be any hardware and software that combine to form the VE.
The model describes thus far only accommodates one individual in the VE but it can be extended to support multiple individuals in a given environment or multiple environments using an external interface, shown as "other environments" to the right of the illustration.

1.5. Defining the Dimensions of Virtual Reality.

Steuer (1992) proposed a system for classifying communication media relative to two distinct dimensions, vividness and interactivity. These dimensions can be used to predict the extent to which mediated interactions will be perceived in a manner similar to their real-world counterparts. Media artist Michael Naimark (1990) refers to these same properties as realness and interactivity. Others, including Laurel (1991) and Rheingold (1991), make similar distinctions. Both vividness and interactivity are likely to have significant effects on social responses to mediated representations. Indeed, many possibilities for future experimentation are suggested by examining the space mapped out by these two dimensions. The next two sections will examine these two dimensions in some detail.

1.5.1. Vividness.

Vividness refers to the ability of a technology to produce a sensory-rich VE. Higher definition leads to a more realistic VR model. There are two important variables that constitute vividness - sensory breath, which refers to a number of sensory dimensions simultaneously presented, and sensory depth, which refers to the resolution within each of these perceptual channels (Figure 1.3.). Sensory breadth is a function of the ability of a communication medium to present information across the senses. Gibson (1966) defines five distinct perceptual systems: the basic orientation system, which is responsible for maintaining body equilibrium; the auditory system; the haptic, or touch, system; the taste-smell system; and the visual system. The vividness of any environment is not generated by any single sensory input alone, but by the simultaneous combination of all sensory input. Traditional media such as print and telephone are relatively low on breath, relying primarily on the visual and auditory channels respectively.

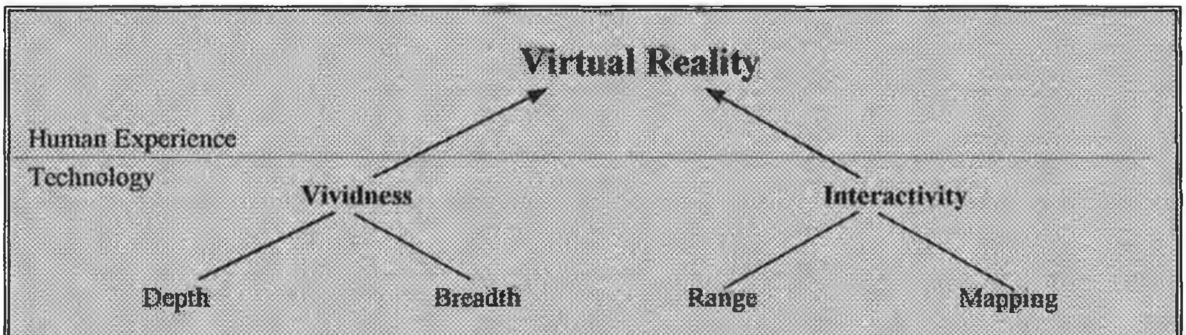

Figure 1.3. Traditional Technological Variables Influencing Virtual Reality.

The vividness of a particular mediated representation also depends upon the depth of the sensory information available in each perceptual channel. This concept can be described in terms of quality: an image with greater depth is generally perceived as being of higher quality than one of lesser depth; the same is true for auditory representation. Informationally, depth depends directly upon the amount of data encoded and the data bandwidth of the transmission channel. The human brain is supreme in its ability to process information; the performance of artificial audio-visual sensors can exceed that of humans but machine interpretation significantly lags human ability. VR is concerned with creating an illusion of reality with the minimum of data. The difficulty lied in creating efficient data patterns that will trigger a desired autonomic perceptual response – given key pattern stimuli, the brain will infer the necessary detail.

1.5.2. Interactivity.

Communication media can also be classified in terms of interactivity. Interactivity refers to the degree to which users of a medium can influence the form or content of an environment. The models produced could have been interactive or non-interactive. Non-interactive models are pre-programmed. The participant views the model and is not provided with a choice at any stage throughout the walkthrough. One cannot decide to turn left here or right there or go straight ahead. In other words, the developer decides the only way in which the model can be viewed. Interactive models, on the other hand, provide the participant with the ability to navigate through the model and to interact with certain objects within the model.

Two factors that contribute to interactivity will be examined here: range, which refers to the number of possibilities for action at any given time; and mapping, which refers to the ability to map its controls to changes in the mediated environment in a natural and predictable manner (Figure 1.3.)

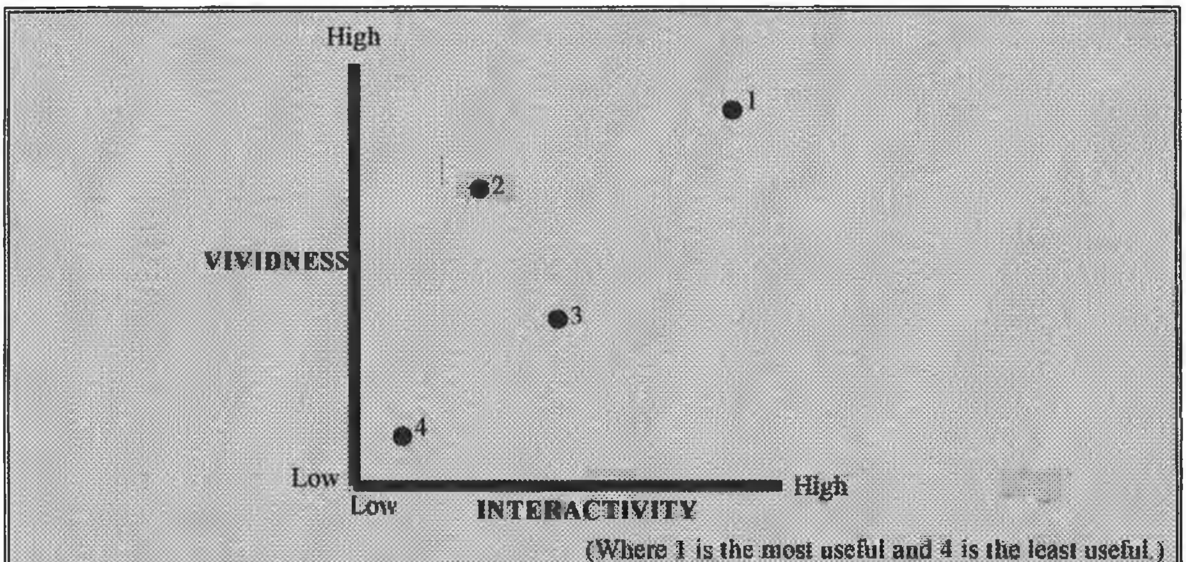

Figure 1.4. Vividness Vs. Interactivity.

The range of interactivity is determined by the number of attributes of the mediated environment that can be manipulated, and the amount of variation available within each attribute. In other words, range refers to the amount of change that can occur within a mediated environment. The specific dimensions that can be modified depend on the characteristic of the particular medium, but include spatial organisation (where objects appear), intensity (loudness of sounds, brightness of images, intensity of smell), frequency characteristics (timbre, colour), and temporal ordering (whether the environment has random access facility or not). The greater the number of parameters that can be modified, the greater the range of interactivity of a given medium.

Mapping refers to the way in which human actions are connected to actions within a VE (Norman 1988). Mapping is, therefore, a function of both the type of controller used to interact with a mediated environment and the way in which the actions of the controller is connected to actions within that environment. At one extreme, mapping can be completely arbitrary and unrelated to the action performed. For instance, an individual nodding might increase the brightness of a sensory output. At the other end of the spectrum, mapping may be completely natural: pointing at an object in a VE might make the person in the environment point accordingly. It is usually advisable to have the mapping as close to the natural motion as possible.

1.5.3. Time.

The two major dimensions across which communication attributes vary, Figure 1.4, have already been discussed, but in order to analyse a VE properly a third dimension, *Time*, must be included, Figure 1.5.

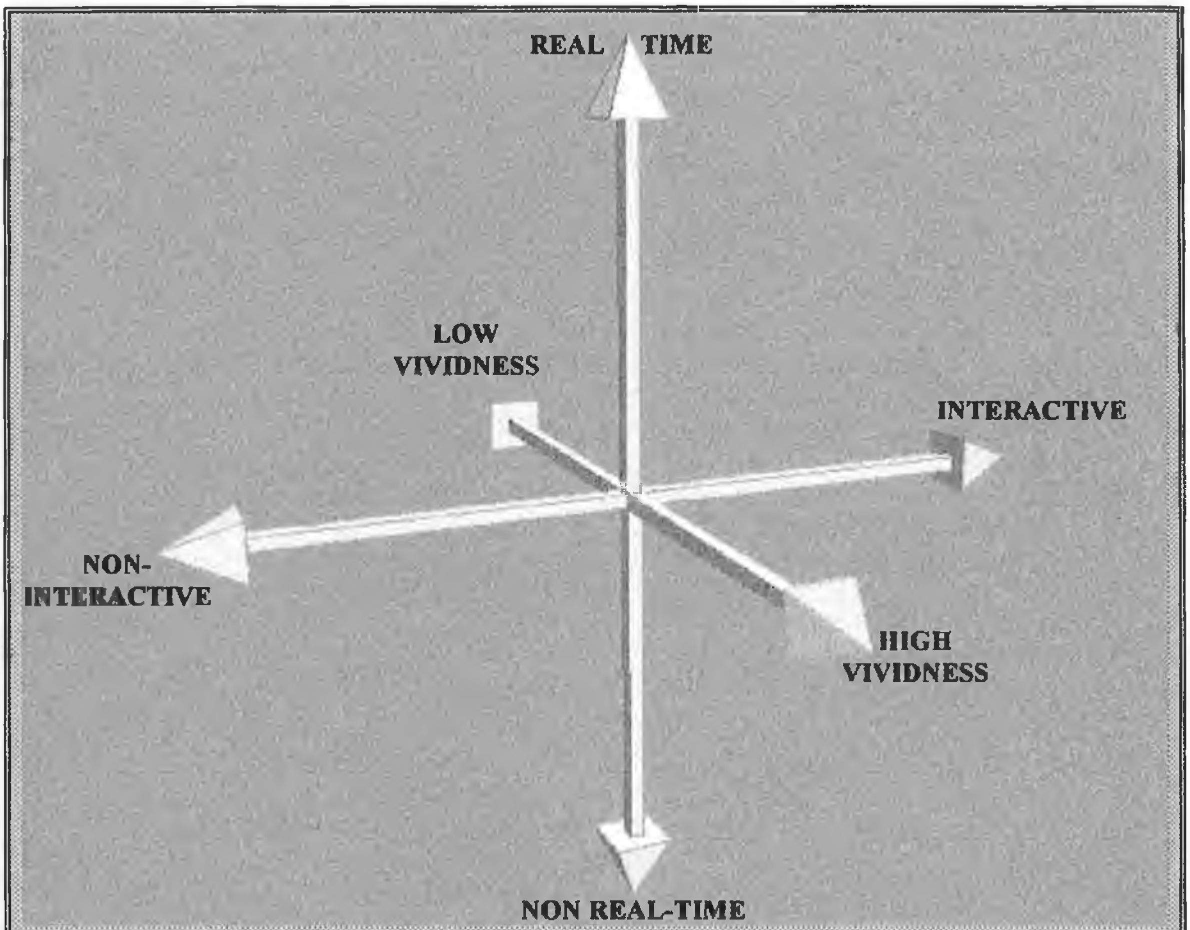

Figure 1.5. The Dimensions of Virtual Reality.

Time refers to the time lag between the instruction and the resulting action; when there is no time lag in the system it is said to be performing actions in 'real-time'. To simplify a little, if the model reacts in the same time as would happen in the real world, then the model is said to be "real-time". If not, a time lag occurs in the application while the computer redraws the scene. This immediacy of response is one of the properties that make even low-resolution applications seem highly realistic. Speed of interaction, or response time, is one important characteristic of an interactive media system (Shneiderman 1992).

There is controversy between researchers at both ends of the "look" versus "walk" spectrum. At one end are those who believe image perfection is the way to reinforce believability. At the other end are those who feel that fluidity of movement is the crucial factor, however elementary the display may look (Brooks 1988). As component technology improves, the gap between the fluidity of movement and the detail will narrow. Ideally, the VE created will both look and behave realistically. It is the aim of this thesis to produce real-time, interactive models of areas of tourist interest with as much detail as is possible. Vividness, real-time and interactivity are not independent. Deciding on one dimension has major implications for the others.

Current VEs are often visually less like the natural environment than many computer games. And VR has no obligation to conform to physical reality. Even when it is possible to create VR systems that are more similar to external reality, part of the attraction of VR may be its ability to construct convincing alternative worlds.

Nonetheless, it may be useful to determine an *ideal* VR as a VE so sophisticated that regardless of how the participant moves or interacts with the environment they will not be able to use sensory cues to determine whether their current environment is real or virtual. The *ideal* VE must, therefore, have perfect vividness, perfect interaction and must occur in real-time. Dennett (1991) makes a strong case that hallucinations indistinguishable from reality are probably impossible. And no simulation can ever be as complex as the phenomenon it models (Rheingold 1991). However, some future VR systems are likely to strive for the ideal.

1.6. Multimedia.

It is appropriate at this stage in the thesis to introduce the second form of technology that is used during the course of this research - Multimedia. The following definition of Multimedia is adopted as it is consistent with the VR definition in that it describes the benefits of multimedia to the user.

> *Multimedia is a computer-based system which allows the integration of various types of information on a computer. In addition to the traditional textual and numerical data, other types of information including graphics, images, audio- and video-clips can be handled in such a multimedia system in digitised form (Vaughan 1994).*

Multimedia is an integration of computing and media. Like VR this union is made possible through improvements in computer technology. Multimedia essentially involves time-based computing. This is most evident when computers are used to author, play and allow interaction with sound and video - the time-based data types. Part of this union is also the expansion of the role of computers as a tool of communications. Thus, multimedia is also the integration of technology into computing in order to provide a more effective communication tool.

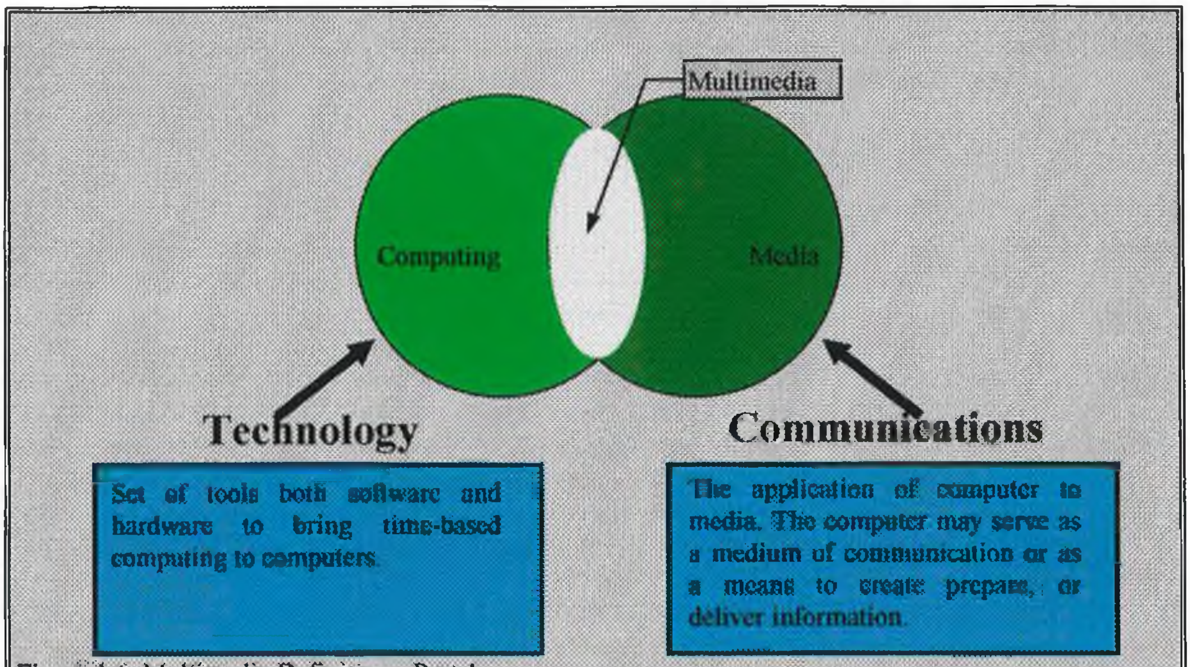

Figure 1.6. Multimedia Definition - Part 1.

When the two components of technology and communication are combined, Figure 1.6. and Figure 1.7., one can see that multimedia is the application of computer technology to spatial and temporal media. It may occur that this is an expansionist view of computers, i.e., why include functions that the computer already performs? In the larger context multimedia is the formalisation of the transformation of computing to become an integral part of media. Viewed in another way, it is an assimilation of media in computers. Desktop publishing was the assimilation of pre-press into the mainstream of computing and the same is happening with media as a whole. One should not infer that media or mass media has become the computer or vice versa. Multimedia, when examined in the context of media, has both inherent strengths and weaknesses, and its development as a market will be moulded by these.

The technical implications of this definition are important. Time-based computing implies that the computers meet some minimum level of performance. For many applications there is a direct trade between quality and the timeliness of a system component or the system as a whole. Multimedia computing means that there must be interfaces to media. Media seldom exists in isolation - the best example being sound and full motion video. Synchronisation is critical to the effectiveness of temporal media. At a system level this impacts principally on data structure and the operating system. When one critically examines the state of multimedia computing, current systems have yet to address all of these design elements effectively.

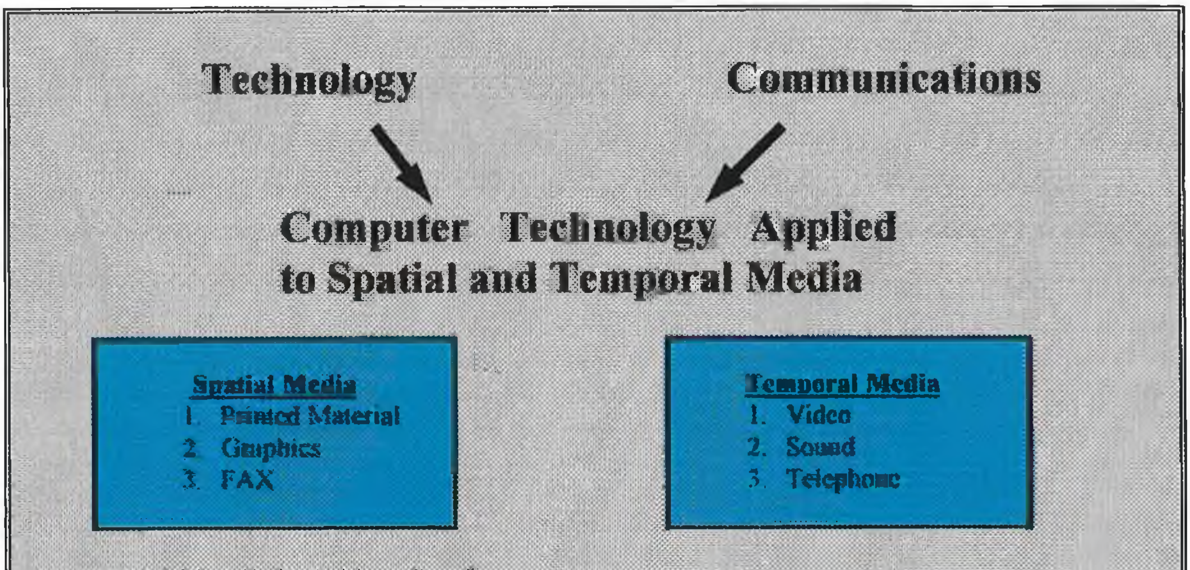

Figure 1.7. Multimedia Definition - Part 2.

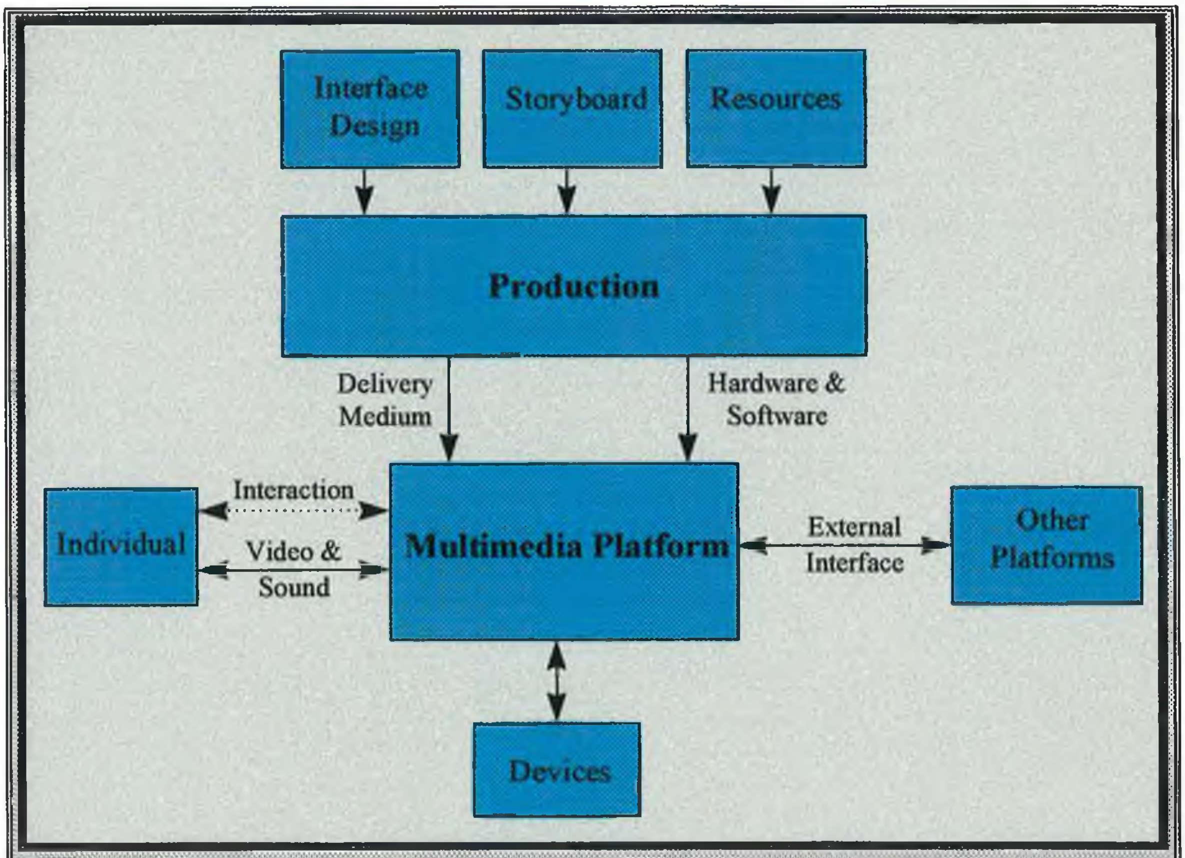

Figure 1.8. Multimedia Model - Non Real-Time.

1.7. Developing a Multimedia Model.

Using the form developed for VR, a corresponding model for non real-time multimedia is developed, illustrated in Figure 1.8. The content of this model is defined in terms of the production process used to create a multimedia application. This process often includes an interface design, the timing of the interaction and a definition of external content such as video and sound sequences. This process is normally conducted with the aid of a storyboard. The final production is assembled and created on a platform. The platform used for production need not be the same as that used for playback. The production may either be resident in the authoring platform or be delivered to a playback platform on a separate medium shown on a line on the model from the production to the platform. Furthermore, multimedia need not be interactive and this is shown as an optional path with the dashed line between the individual and the multimedia platform. The most frequently used analogy for this form of multimedia is book publishing, especially when the delivery medium is CD-ROM. Yet, when one examines the definition and the potential of multimedia, such an analogy is quite limiting. When an analogue medium is used it is attached to the platform as shown in the lower portion of the illustration. Multimedia can also be linked to other media or other platforms.

The model for real-time multimedia, Figure 1.9., is a decided shift from the non real-time version. The key reason is that the system can operate at any time and may run continuously. This model includes a design for the environment and a model for how operations are conducted. There are close analogies with both television station operation and computer network operations. Multimedia, in this environment can run continuously and change dynamically based on content and the action of participants. As such, one defines an "operations environment" which forms the basis for real-time operations. This model also places a greater emphasis on interface to media sources. Interaction is optional as in the previous model. When interaction is not used this form of multimedia is clearly passive. Such an approach neglects some of the advantages of multimedia computing.

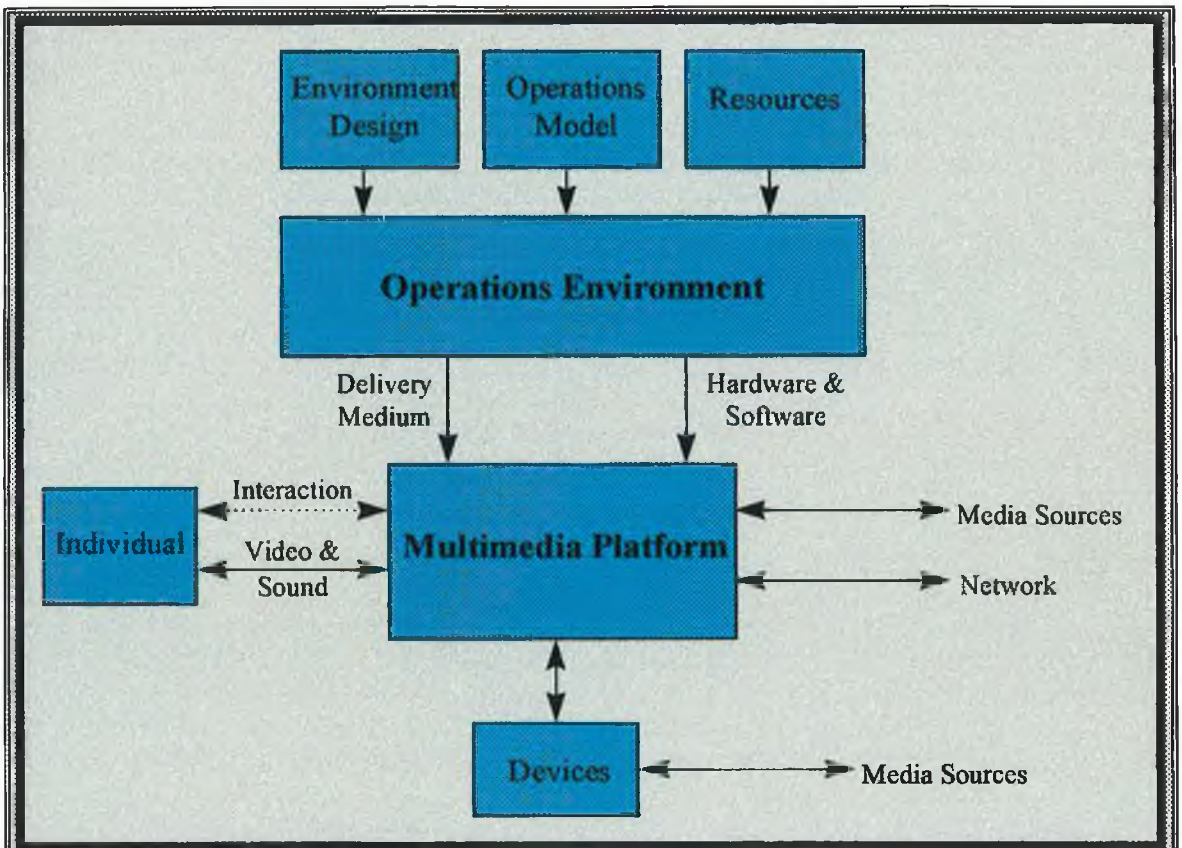

Figure 1.9. Multimedia Model - Real-Time.

There is a clear pattern between the three models: VR, Figure 1.2, non real-time, Figure 1.8, and real-time multimedia, Figure 1.9. Each is quite different, yet, there are common elements and technologies required to implement them. It is also instructive to compare both the VR model and the real-time multimedia model. Many of the components are the same. A major difference is in the role of the user interface in VR.

The synergy of VR and Multimedia is illustrated clearly in their union in Figure 1.10. They share the need for system performance and the integration of time throughout each. Multimedia is distinct in its interface to media and the input and output (I/O) devices it uses. Although it can be argued that VR also uses displays and speakers, VR is focused on the human interface, the appliances used by the participant and the definition of the VE. One can also expect to see increasing convergence of these two technologies in the future.

1.8. Summary.

The focus of this research is on the development of realistic, real-time and interactive VR models of areas of tourism interest. The first step to be taken is to provide a comprehensive understanding of the concepts of VR. This process began by choosing an appropriate definition from the three main approaches to defining VR. These three approaches are in terms of the technology used, in terms of presence, and in terms of the VE created. After the different types of VR definition are discussed the definition which is most appropriate to this research has to be decided upon. This definition is the definition in terms of the VE created.

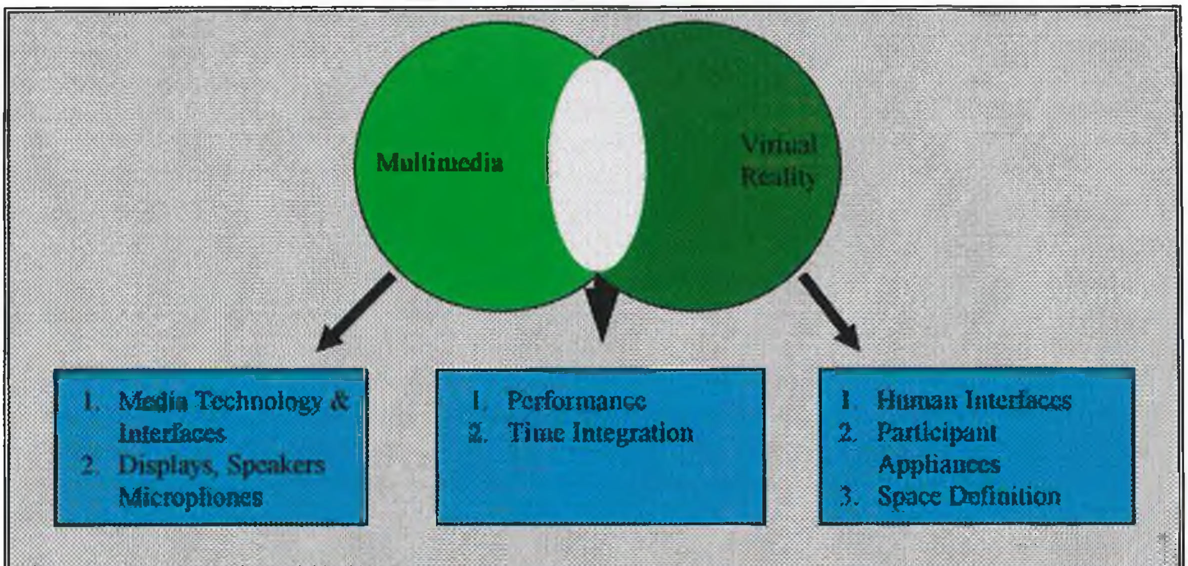

Figure 1.10. The Union of Technology.

Once a definition is decided upon the next issue addressed is an examination of the evolution of VR. VR is an integrative technology - consisting of many different, and often, unrelated technologies. Therefore, the development of VR is dependent on the development of its component parts. In order to trace the evolution of VR, the evolution of its component technologies has also got to be traced. This is achieved by the use of a two-tier approach. The first tier traces the conceptual evolution of VR to its origins in 3D cinematography in the early 1950's and the second tier traces VR's physical evolution through the development of its component technologies. VR, as it exists at present, has occurred only in the last ten years due to the critical convergence of technology. A development toward greater sensory vividness and involvement on the behalf of the individual may be observed. More unexpected perhaps is the technological realisation of the desire to reach out and touch and interact with the media images that are created. What truly sets VR apart from other media is the technological imperative to directly put the body of the user inside the illusion, to surround the user with a space that stretches infinitely in all directions, a world of unexpected experiences. With appropriate interfacing equipment, users can visit, navigate, and interact within the interiors of these 3D environments.

A comprehensive VR model suitable for the classification of any VR system is then developed. This model is designed to describe the basic generic framework in which any VR system may be implemented and, thus, every system may not contain every element of this model.

The dimensions of VR are the next element of VR to be explained. Traditionally, communication media are classified relative to two distinct dimensions; vividness and interactivity. In the case of VR these have to be widened to encompass a third dimension, time. It is along these three dimension that each VE can be classified and subsequently placed on a 3D graph in accordance with their classification. An ideal VE is one in which a participant can not distinguish whether it is real or virtual. An ideal VE is one with perfect interactivity, perfect vividness and respondents to a participants movement in real-time.

The final part of this chapter examines the concepts of multimedia. In order to explain these concepts fully a definition of multimedia is provided and both a non real-time model and a real-time model were developed. Certain similarities between the VR model and the real-time multimedia model are evident. The main reason for these similarities is due to the fact that both models share common elements and technologies. The final section of this chapter deals with the synergy of VR and multimedia and the possible increased convergence of these two technologies in the future. The following chapter deals with the types of VR systems currently available and their hardware and software components.

Chapter Two

Virtual Reality - The Systems And Components

2.1. Introduction.

The main purpose of this chapter is to examine the different types of VR systems available in order to choose the type most appropriate to the tourism industry. There are currently three main types of VR system available. They are Immersive VR systems, Third Person VR systems and Desktop VR systems. A VR system is actually comprised of an array of possible input and output devices, each serving a sensory channel or linked to the participant's body movements and responses. These independent sensory inputs and outputs have been developed to produce the kinds of visual, audio and tactile effects used in current VR systems.

Each of these sensory inputs re-inforce one aspect of the participant's illusion of immersion within the Virtual World. All of these inputs not only draw from and are integrated with the available hardware and software setups, but also act in tandem with one another to create this illusion. In an attempt to explain the three types of VR systems currently in existence one must first examine the hardware and software components which converge in order to make these systems possible, Figure 2.1.

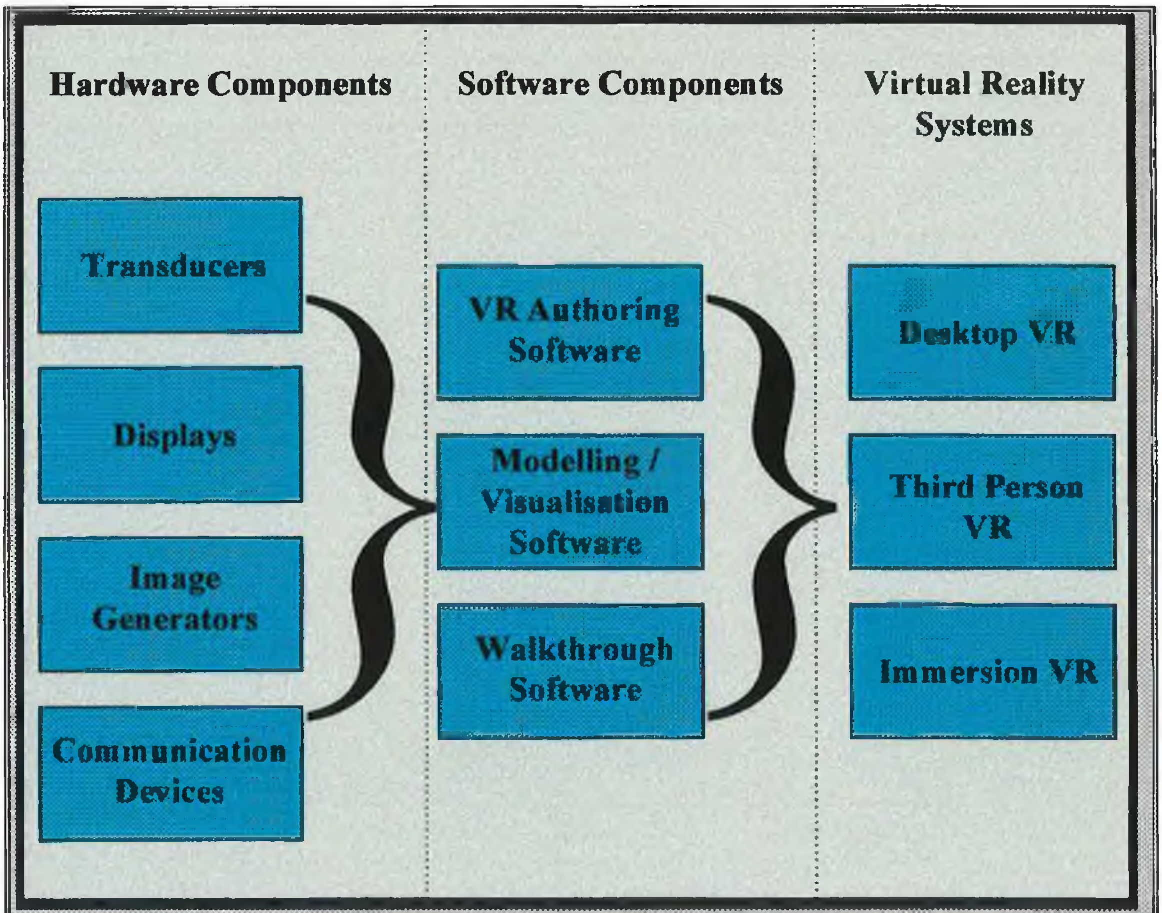

Figure 2.1. The Components of Virtual Reality Systems.

2.2. Virtual Reality Hardware Components.

Interface devices are the translators between a participant - their perceptual and nervous systems, their minds and bodies - and the computer model. Without these devices, a participant is excluded from the Virtual World. A major distinction between VR and other forms of computer simulation lies in what is called “inclusiveness”: the ability of the participant to interact with the computer-generated environment as though he or she were actually inside a wholly contained world (Bricken 1991).

There are two classes of device required to make a simulation inclusive. The first class consists of transducers that bring the user into the simulation by tracking head, hand and body motion and by sensing intentional actions and the second class are displays which in turn bring the simulation to the user. Displays not only refer to visual displays but also auditory and tactile information presented to the user in such a way that it replaces the corresponding senses arising from the physical world. The design and choice of the devices that perform these tasks shape much of the VR experience and the devices appropriate in given situations depends on the purpose of the particular world.

In VR, as perhaps in no other medium, the user interface designer is confronted with the task of designing tools that are consistent with Virtual Environments (VEs). Although present VEs may only simulate a few sensory modalities: vision, sound and perhaps tactile feedback, people always take their entire nervous systems with them. VR can never be passive, because visual displays in particular must always match an individual’s own motions. This intimate association with the user’s perceptual system is also the reason that response time is critical in the design of all components of VR systems. A lag between action and response that may be perfectly acceptable in many computer applications can be very unsettling in a Virtual Environment (VE).

The following section describes the hardware components necessary to translate a participant’s commands, queries and responses into computer activity in order to deliver and display a VE, to store the Environment, and to interact with other participants while inside the Environment. In order for this scenario to be achieved there are four components to be considered:

1. Transducers,
2. Image Generators,
3. Display Devices,
4. Communication Devices.

The distinction between these four components is not always clear because a good VR system has to incorporate all of them seamlessly. Therefore, even though each component will be discussed separately it must be understood that there is significant interdependency.

2.2.1. Transducers.

In order for a person to interact with a VE, their actions need to be communicated to the reality engine. This process is carried out by a device known as a transducer. Early input devices were cards punched with series of holes that were meaningful to a binary machine, the computer. A later, and still current, input device is the keyboard on which we type the data into the computer. Human frustration with these devices led to the development of other, more “user-friendly” devices for inputting data. It is the responsibility of the input device to get the co-ordinate data to the rest of the system with minimal time lag. The more intuitive the input device, the more potential there is for interaction between the computer and the participant and among participants. Ideally a transducer should be perfectly intuitive, real-time, and should also provide a participant's position normally expressed in the Cartesian dimensions of X, Y, and Z, and orientation expressed in terms of roll, pitch, and yaw. These six directions are known as the six degrees of freedom (6DOF) and are clearly illustrated in Figure 2.2.

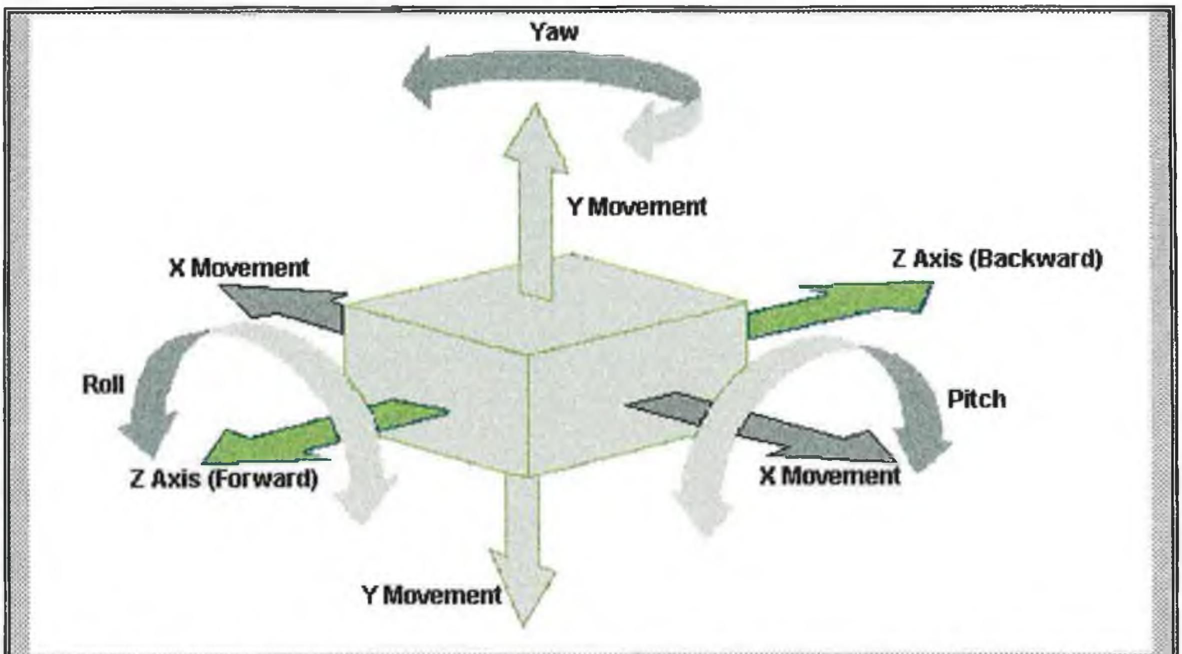

Figure 2.2. Position Measurement (X,Y,Z axes) and Orientation Measurements (Pitch, Roll and Yaw).

The role of the transducer is to convert an action into a form that can be interpreted by a computer. These actions include movement, speech and brain activity. Each of these actions shall now be dealt with individually.

2.2.1.1. Movement Transducers.

Movement transduction is necessary for a person to move naturally through the VE. All movement transducers can be placed into one of the following five categories: controllers, position trackers, gloves, bodysuits and retinal scanners. In this section these five types of transducers are examined.

Controller Devices: The simplest controller devices used in Virtual Worlds are the conventional *mouse, keyboard* or *joystick*. While these are two dimensional (2D) devices, creative programming can use them for six dimensional (6D) controls. There are also a number of 3D and 6D mice and joysticks currently on the market. The principle of each device is quite similar, namely to produce multi-axis control inputs into a computer for the manipulation of virtual objects. They achieve this by adding some extra buttons and wheels that are used to control not just the X-Y translation of a cursor, but its Z dimension and rotations in all three directions.

Trackballs are 6D pressure-sensitive input devices that can be used to select and manipulate objects or control a participant's viewpoint. They are similar to mice, except that the cursor is moved using one or more fingers to roll across the top of a ball. Moving a viewpoint through a Virtual World with a trackball is very smooth. Trackballs have at least two buttons; one for the user to click or double click, and the other to provide a press and hold condition necessary to select from menus and drag objects. Some trackballs have more than the mandatory two buttons, as is the case the trackball in Figure 2.3.A. These additional buttons may be pre-programmed to carry out specific tasks by the user. This option can be very helpful and time saving when used effectively in the design process. Unlike a glove, it can be used to tumble and swivel. This capacity is very useful for object manipulation, but moving a viewpoint with unconstrained freedom can sometimes be disorienting.

Figure 2.3.A. The Trackball   Figure 2.3.B. Touchscreen Monitor.

Touchscreens, Figure 2.3.B., are monitors that are sensitive to pressure and registers the location of the user's finger when it touches the screen. The mechanics of touchscreen technology vary but the underlying principle is the same. Most touchsreens have a textured coating across the glass face. This coating allows the user's finger position on the screen to be recorded and translated into a form that the computer can understand. Other touchscreens use beams of infra-red light that criss-cross the front of the monitor to calculate where exactly the finger has been pressed. The more expensive and complex touchscreens technology not only measure the location where the finger was applied and the force that was exerted on that location, but also determine the user's position and orientation. A keyboard is sometimes simulated using an on-screen representation so users can input text by simply pressing the keys. When an application is designed to use touchscreen technology the monitor is the only input device necessary. Touchscreens are not recommended for day-to-day computer work, but are excellent for multimedia applications in a kiosk, at a trade show, or at an interpretative centre.

The Wand is an another evolution of the Mouse. It is a simple physical device with a wide diversity of uses, ideal characteristics for a tool. Physically, the wand is a spatial position and orientation sensor on a hand-held stick. Most incorporate on-off buttons to control variables in a simulation or in the display of data. Others have knobs, dials, or joysticks. Their design and manner of response are tailored to the application. For example, biologists sometimes use wands like scalpels to slice tissue samples from virtual brains. Most wands operate with 6DOF. This versatility coupled with simplicity are the reasons for the wand's popularity.

Position Trackers: One of the key components of the VE simulation system is the means provided for the participant to interact with this environment. This requires tracking the position of a real world object, such as a eye, head, hand or body, and using this data to control the movement of a virtual representation. These movements are then translated into position and orientation coordinates, which are in turn deciphered by the computer. Images corresponding to the viewpoints can then be displayed. There are four elements required for position tracking:

Accuracy:	The fractional error in making a measurement.
Resolution:	The smallest movement that the instrument will detect.
Latency:	In the case of position trackers, latency is the time between when the measured object is located and when this position and orientation data set is available in a tracker output register.
Speed:	Speed, in this context, is the maximum number of complete 6DOF updates per second that the instrument will produce.

For tracking purposes, a physical device is attached to the real world object so that the participant's movements can be detected. Movement has been measured electro-magnetically, optically, ultrasonically, gyroscopically and mechanically or with combinations of these methods. Each of these systems is described in more detail in Appendix A.

All current types of position trackers put constraints on the user. Many tracker systems require a clear line of sight at all times between the sensor and the target of the user. Drawbacks to ultrasonic sensors include low resolution, long lag times and interference from echoes and other noises in the environment. Optical tracking systems such as the University of North Carolina (UNC) ceiling tracker (Wang 1990) promise higher resolution range, and sample rates, with reduced latency. The challenge is to improve the technology for wider range, wireless communication, faster update rate, and greater accuracy.

Whichever tracking system is used, the principle is the same; the computer tracks the movement of the object, and alters the position and orientation of the object in the VE accordingly. For head tracking, sensors are placed on top of the head-mounted display (HMD). As the participant's head moves, calculations are conducted, and the image is regenerated and updated to coincide with the head movement. These images reflect changes in a VE corresponding to what the user would expect to see if the movement was made in a real world situation. Sensors can be mounted elsewhere such as on gloves or bodysuits.

Gloves: Probably the most ubiquitous device for control and input to a VR system is VPL's DataGlove. The DataGlove consists of flexible optical fibres running along each of the hand joints. Each finger and thumb of the lycra glove has a loop of optical fibre cable running up and down its length, with an light emitting diode (LED) at one end and a photo-transistor, a device used to convert light into an electrical signal, at the other. When a finger or thumb is moved, the amount of light transmitted changes. On the back of the glove there is an electro-magnetic sensor which tracks the hand's overall movements. On screen, an image of the hand moves in real-time, shadowing the movement of the hand in the DataGlove and immediately replicates even the most subtle actions. DataGloves are separate devices that can be interfaced with any computer system, and most applications, from video games up through multi-million pound military research projects. The Mattel PowerGlove, Figure 2.4., uses ultrasonic transmitters and receivers and a triangulation technique to achieve the same purpose. The PowerGlove was originally manufactured for use with the Nintendo game system. This device is easily adapted to interface to a personal computer (PC).

Figure 2.4. The Mattel PowerGlove.

The use of a glove input device increases a participant's sense of presence and allows him to manipulate virtual objects in a more natural, efficient and intuitive way (Sturman et al. 1989). However, gesture commands are limited by the number of finger positions that can be defined as unique, mutually exclusive, and comfortable. When moving by pointing, directional accuracy is approximate, stopping exactly where you want to can be difficult, and speed control is awkward. Inadvertent gestures can trigger unintentional commands. Gestures can also require a degree of manual dexterity that many people do not have. There are also proprietary issues regarding the use of computerised gloves that will affect price, performance and availability.

Bodysuits: The bodysuit is basically a customised DataGlove for the whole body. It is instrumented with a fiber optic cable running up and down all the major body joints. As a user moves, bends or waves, the system collects spatial coordinates from each part of the bodysuit and dynamically updates a full range of body actions in the VE. The movements of a bodysuit are often restricted by the wire connection. Walking, for example through the design of building, has been simulated by a navigable treadmill. A fixed bicycle, rowing boat or car would provide similar effects for appropriate environments. Currently, large movements through the VE are commonly achieved by having the person fly, using hand motions to control the direction and speed of travel.

Retinal Scanners: Eye-tracking technology, like sensor technology, has many approaches. Eye movement is probably more important as a control input than as a prerequisite for achieving an effective simulation. The state of the art for eye movement transducers is quite advanced, having been driven by medical/psychological research. Options include transducing electrical activity around the eye, corneal reflection and direct imaging followed by video image processing. Most systems are mounted on helmets or eyeglasses, but some are remote. A variety of commercial systems has been widely used, mostly by researchers in perceptual psychology. Nearly all of the modern systems use video cameras to detect light from infra-red LED's irradiating the cornea. A number of these systems have as their aim the creation of a large field of view, high resolution display by using the fact that vision is most acute at the optical center of the retina.

2.2.1.2. Speech Transducers.

Voice recognition systems have been under development for a number of years, driven by potential commercial applications such as word processing, military applications, and control input applications. Although far from a mature technology, existing systems are capable of recognizing a limited vocabulary and are being marketed for the home computer market. Voice recognition systems facilitate hands-free interaction with a VE. These systems usually provide a unidirectional, noise-canceling microphone that automatically filters out background noise. Most types of voice recognition systems currently available can trigger common menu events such as Save, Open, Quit and Print. In the case of VR, voice recognition systems must first of all be taught to recognise individual voice commands, such as Forward, Backward, Right, Left, Walk and Stop, and then be programmed with the appropriate responses to the recognised commands. Voice recognition systems in VR may be another tool for increasing the intuitiveness with which a person interacts with a VE.

2.2.1.3. Thought Transducers.

Although bordering on the realm of science fiction, some research has been conducted on the possibility of using thoughts, actually measurable brain activity, for controlling movement within a VE. The potential use of electrical brain activity as a control input is limited by the lag imposed by the requisite filtering of the raw signals. The Air Force are interested in this area because of the potential for monitoring, and having systems respond to, the psychological condition of their pilots while they are engaged in high-stress activities.

2.2.2. Image Generation.

One of the biggest problems that face researchers in the area of VR hardware is in the field of real-time image generation. Visual image generation is just one part of what has to be achieved. Artificial stimulation to the other senses must also be considered. However, since the image of something in a particular sensory modality is only as useful as the ability to display to that modality, most effort has been expended on visual and auditory images. In this section both visual and audio image generation along with haptic image generation are discussed.

2.2.2.1. Visual Image Generation.

In VR systems, the requirements for real-time response dominates all others, and one of the most onerous tasks is to generate sufficiently realistic images at acceptable frame rates. Therefore, the processing power machine used to produce VEs is very important. Most high quality processors can generate sharp, complex images slowly or they can produce simple, fuzzy images much faster. In the case of an ideal VE a photo-realistic image must be generated in real-time. This scenario has an almost insatiable appetite for a computer's processing power, and the failure of most systems to generate such high-quality images at such high speed may prove to be a major stumbling block to the development of realistic VEs. However, there are certain hardware devices available, such as 3D accelerator cards, which assist the processor to render high quality VE in real-time. Therefore, not only is the speed of the processor used to produce the VE important, but so too is the 3D accelerator card that the machine possesses. With this in mind, one can expect the next major development in the field of VR to be 3D accelerator cards.

2.2.2.2. Audio Image Generation.

Auditory systems are very important as an alternative or supplementary information channel. Creating some level of audio environment can greatly enforce the information presented in the visual display. This information may be natural, such as familiar sounds associated with familiar objects, or completely unnatural, such as proximity sounds, collision detection, or general warnings that increase our ability to assimilate and manipulate information within the VE.

The value of auditory systems can be further enhanced by providing a truly 3D auditory display in which sounds are actually localised in space relative to the participant. Sounds in the real world are subconsciously localised and used by people, for perceiving warnings, for orientation and for distinguishing conversations against a noisy background, such as the cocktail party phenomenon. Several components of a sound are used to determine its spatial origin. The time difference between a sound reaching each of our ears, and the amplitude of the sound in each ear, are important cues for determining its origin. This approach is particularly relevant in the case where spacial awareness is important, e.g. when visual cues are limited or absent.

2.2.2.3. Haptic Image Generation.

The haptic senses are the pressure senses and the muscle mediated senses that tell one where one's limbs are in space. In a haptic display objects seem to push and pull in different direction with varying degrees of strength and resistance. Objects in a VE may be assigned mechanical properties such as force, torque, friction, heat and pressure which becomes noticeable to a participant interacting with that object. The value of such senses in both basic manipulative tasks and in more complex visualisation problems is extremely important (Brooks et al. 1990). However, these systems currently have not advanced significantly beyond the research domain, for the primary reason that they are extremely expensive and have severe ergonomic problems.

Signals from a haptic display are meant to be interpreted by a user so that they can adjust movement accordingly in order to achieve some goal. Participants in particular VEs become accustomed to certain cues and use them to estimate real depth or force or distance. Practice improves both performance and the level of perceived realism. Force feedback, even from the most simple of virtual objects, is a very difficult task, and haptic displays are not designed as "touch" machines but rather as environments from which a person can gain information associated with the properties represented, such as weight or solidity, from the signal delivered through the haptic display. Despite the quality and intensity of the feedback one can not currently sit on a virtual chair. Gloves and bodysuits are gradually being lined with areas that provide haptic effects to their wearer. These are connected to the VR system and controlled to add realistic effects to the VE.

2.2.3. Display Devices.

The effect that a VE has on a participant depends substantially on the way the environment is viewed, heard and possibly felt by that participant. In order to experience this realistic VE the participant is very much dependent on the visual, audio, and possible haptic cues that they receive from the VR system. There are a range of display systems which have proved essential in creating realistic VR experiences. In this section these Visual, Audio and Haptic display systems will be discussed.

2.2.3.1. Visual Displays.

Sight is our most important sense (Lovine 1995) and this is the main reason why most VR research attention is focused on visual displays. There are three key areas that must be taken into consideration when dealing with visual display devices. They are:

1. Graphics Resolution,
2. Field of View,
3. Stereo Imaging.

Graphics Resolution refers to how many pixels are used in the display. A standard VGA screen is 640 pixels wide and 480 pixels high. This space is very restricting. What is even more restrictive is that a standard VGA screen only support 16 different colours. These two constraints make it extremely difficult to create a realistic VE. There is, however, another side to the VGA story, and that is a video mode that is 320 pixels wide by 200 pixels high but allows for 256 colours. This video mode is simply a trade-off - trading the number of pixels for the number of colours. Giving up pixels, of course means giving up resolution but many early VR products support this low resolution mode because any PC with a VGA card can support it. One does not have to use the standard VGA resolution; one can use the Super VGA or a different type of adapter setting to achieve a much higher resolution and a larger number of colours.

The main point that must be taken into consideration when dealing with resolution is that the more pixels that a display uses, the more processing power is required to generate the images. In this situation the key component is the video card. The faster it operates the smoother the transition in the VE will appear.

Field of View (FOV) is normally thought of as an angular measure of visual perception for a given head position. One of the hardest things to achieve in a VE is a wide FOV. The width of a FOV is often taken for granted but it will be quite some time before affordable VR systems offer adequate displays to fill a person's peripheral vision. A computer monitor has a very narrow FOV as it occupies only a small proportion of a person's FOV. To achieve a wider FOV one can get a larger monitor. The average monitor measures 14 or 15 inches diagonally, one could upgrade to a 17 or even a 21 inch monitor. This unfortunately, will not substantially increase the FOV. To get a wider FOV a person needs to build a wall size negatively curved monitor, use a projection system, or use a HMD.

Stereo Imaging refers to viewing an object in duplicate to create an illusion of 3D. Stereo imaging is accomplished by creating two different images of the world, one for each eye. The images are computed with the viewpoints offset by the equivalent distance between the eyes. There is a large number of technologies that can present displays in stereo vision but one of the main problems associated with stereoscopic displays is the conflict between convergence and accommodation. When one looks at a close object, one's eyes turn in, this is known as *convergence*, and when one shifts focus from a distant object to a near object, one's eye muscles adjust the shape of the lens to focus at the proper distance, this is referred to as *accommodation*. When the focal distance is at odds with the convergence angle, some observers complain of headaches or of not being able to merge the two images into a stereoscopic image. Developers are working now on decreasing the effects of convergence and accommodation conflicts.

Computer Monitor: VEs can, of course, be experienced on a commercially available PC with no additional hardware. In such a case the monitor will suffice as a display device. A wide variety of monitors is available and they vary in both the video modes that they support and in their size. High-end large screen, negatively curved graphics monitors are available but they are expensive. A point that must be taken into consideration is that an application should be developed on a monitor of the same size and resolution as the monitor that will be used for the application's delivery.

Head Mounted Displays: A HMD uses some form of helmet or goggles to place small video displays in front of each eye, with special optics to focus and stretch the perceived FOV. Most HMDs use two displays and can provide stereoscopic imaging. Others use a single larger display to provide higher resolution, but without the stereoscopic vision. The HMD is currently the most important part of the arsenal of equipment for a visitor to an Immersive VE. It is also one of the most complex pieces of equipment used in VR today. A HMD provides greater verisimilitude to an image by adding two perceptual cues to the standard video image, stereo imaging and head-coupled motion parallax cues. What this means is that when the participant moves his head, the stereoscopic image changes just as it would if the viewer were looking at a real world scene.

Figure 2.5.A. Liquid Image's Opaque Display.

Figure 2.5.B. Virtual Vision's Transparent Display.

The are currently two main types of HMD available. They are opaque HMDs and transparent HMDs illustrated in Figure 2.5.A and Figure 2.5.B. respectively. An opaque display is one in which the view of the VE blocks out and supplants the view of the real world. A transparent display, on the other hand, is one in which the virtual images seem superimposed on the real world, which remains visible through the visor. In effect, an opaque display replaces the real world. When a user chooses transparent display over opaque, he continues to rely on the real world as a frame of reference. This is more "augmented reality" than "virtual reality".

All HMDs are based on aerial imaging technology, in which we view a picture focused at infinity that is projected in front of our eyes. This, of course, is not the way we see objects in the real world, where our focus changes to converge on objects located at different distances from us. Although binocular viewing does give us a feeling of three-dimensionality, it is difficult to judge the relative location of objects in aerial displays without redundant cues such as familiar size, occlusion and parallax (Furnass 1990).

In summary, most current HMD's are bulky, heavy and fairly fragile. Although things are improving most displays are very low-resolution, making small objects and details of Virtual Worlds difficult to see. All HMDs have cables that restrict our movement and some distort vision with a fish-eye effect, and others do not show colour. Most do not include audio capacity.

Binocular Omni-Orientation Monitor 2: Binocular Omni-Orientation Monitor 2 (BOOM2), illustrated in Figure 2.6., is a mechanically supported stereo-scopic viewer. It provides the user with a high resolution, physically comfortable alternative to currently available HMDs. As a floor standing device, the boom requires little recalibration and customisation and can be passed easily among a number of users. It is comprised of a long rod perched at its center, swings around upon a jointed frame with 6DOF movement. Optical encoders at each of the six joints of the BOOM2 feeds the system this high-precision information about the position of the participant's head. The viewing box is suspended from a swivel frame at one end and weight counterbalanced by a weight at the other. Two handles are located near the base of the box and are used for pulling and holding the box to the user's face and navigating it to display different views of the VE. The BOOM2 is operated much like a pair of binoculars, except that the counterbalanced linkage keeps the view steady, requiring only one hand for use. Thus, the free hand can easily manipulate tracking devices or other equipment. The hand controlling the BOOM2 can also simultaneously operate additional instrumentation, enabling the user to manipulate virtual interfaces.

Current models present monochrome, colour and "pseudo" colour scenes on a cylindrical display space about five feet in diameter and five feet high. Its FOV, at any one time, is about 140 degrees horizontally by 90 degrees vertically. Its stereo effect is basically achieved by rendering red-shaded images to the left eye and blue-shaded images to the right eye. It is designed for the real world, as a complementary tool to the computer, set up at the desk to enable the user to quickly look into the world just designed. Group interaction is also simple, because the BOOM2 can be passed effortlessly from user to user. The BOOM2 built-in mechanical tracking devices also eliminates the need for an electro-magnetic tracking system, avoiding the problems associated with electro-magnetic devices described in Appendix A.

Figure 2.6.Binocular Omni-Orientation Monitor (BOOM).

The use of a BOOM2 has certain advantages over the use of a HMD, especially if it is to be used for a long period of time. Because a HMD is actually worn by the participant there are limits to how heavy it can be or how long it can be worn before becoming uncomfortable. Most HMDs are fitted with lightweight LCDs, which produce only grainy, low resolution images. A BOOM is not worn nor is the weight borne by the participant and, therefore, they can make use of CRT technology capable of producing sharp, detailed, high resolution images.

Red/Blue Glasses: This type of display technique is the easiest form of 3D technology, but is also the least effective. To view the 3D image the observer must wear glasses that use a blue film in one eye and a red film in the other. The source image, normally viewed on a conventional computer monitor, contains overlapping blue and red versions of the scene, illustrated in Figure 2.7. To create this image a stereo-pair is needed. This is a term used to describe the two images that represent what each eye would see if it were substituted for the camera position in the scene. In this scenario, one image is encoded using blue colour only, and the other image is encoded using only red colours. When the image is viewed the red image is attracted to the right eye, and the blue image to the left eye, due to the red/blue films in the glasses, resulting in a basic 3D image.

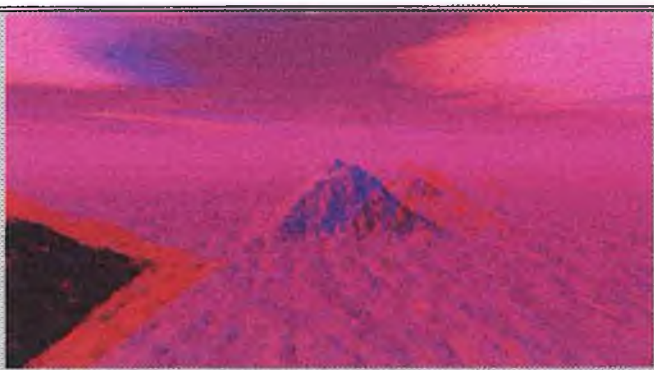

Figure 2.7. Red/Blue Image Display.

Shutter Glasses: Shutter glasses comprise two lens each containing a shutter that simply blocks the view on that side when the eye is closed. This method works well with video sources as a video is actually made up of two fields. In the case of VR the image for one eye is put in one field and the image for the other eye is put in a second field. The two images are then displayed sequentially on a conventional monitor or projection display. Liquid crystal shutter glasses are used to shut off alternate eyes in synchronisation with the display. When the brain receives the images in rapid enough succession, it fuses the images into a single scene and perceives depth. A fairly high display swapping rate is required to avoid perceived flicker.

The main disadvantage associated with shutter glasses is that they effectively cut the vertical resolution in half. This disadvantage is offset by the fact that there is little or no special hardware requirements in order to encode the images. However, locating the glasses is difficult as few are still being made or sold for their original use. An example of a shutter glass system is illustrated in Figure 2.8.

Figure 2.8. Stereographics "CrystalEyes" Shutter Glass System.

Figure 2.9 The Cyberscope.

Cyberscope: Another alternative method for creating stereo imagery on a computer is to use one of several split screen methods. These divide the monitor into two parts and display left and right images at the same time. One method places the images side by side and conventionally oriented. It may not use the full screen or may otherwise alter the normal display aspect ratio. A special hood viewer, similar to the one illustrated in Figure 2.9., is placed against the monitor which helps the position the eyes correctly and may contain a divider so each eye sees only the appropriate image.

Figure 2.10. The Mandala System - An Example of Projected Displays.

Projected Image: This type of display requires no HMD, gloves or suits; instead a video system merges a real-time image of the user with a computer generated scene which is then projected as a distant image which fills the participant's FOV. With projection systems, since the display device is not worn, fatigue is low. Maintenance costs are also generally lower due to less need for instrument recalibration. The system can also be linked up to a computer screen with only minimal image distortion occurring. Participants integrate their movements with the display and appear to interact with each other and with the virtual images. This is the technique that is used in the Mandala system, a Third Person VR application, illustrated in Figure 2.10.

Other Approaches: There are some exciting new approaches to 3D displays without goggles. Sanyo's 3D display uses two LCD projectors behind a double lenticular, ridged screen. The images are based on signals from eight separate cameras, organized in four pairs, which are interleaved in vertical columns and projected through the screen. The lenses on each side of the screen bend the light so that you only see the image from one camera at a time with one eye. The other eye sees the image from the corresponding camera.

Another approach under development at Terumo Corporation in Kanagawa, Japan, uses a pair of LCD panels along with monochrome CRTs, or LCDs, as backlights. A half-mirror optically combines the two images. By controlling which portion of the CRTs are lit, the light can be directed to just the left or right eye. For this to work, however, lighting must be located precisely in relation to the observer's head. The display includes a sophisticated infra-red tracking system that locates the observer's face, and, using a computer, calculates the center of gravity for that location.

Dimension Technologies of New York, has a no-goggles 3D display that vertically interleaves a pair of images from separate cameras. The backlight is focused into thin vertical lines. Light traveling from one backlight strip will reach one eye after traveling through a column of even-numbered pixels and the other eye after passing through odd-numbered pixels.

2.2.3.2. Audio Displays.

Sounds can be used to supplement visual information that is becoming too cluttered. They can also serve as a reminder to conduct a certain action. In addition, sound information within a VE gives a participant an exact location of a sound source e.g. running water or music. With the multimedia speakers that flank most monitors today, stereophonic audio is dispersed in a flat plane between the left and right speakers. With 3D audio, those same speakers can make sounds that seem to come from over your shoulder, which when combined with realistic images can result in truly realistic VEs.

There are several competing technologies for 3D audio. The Sound Retrieval System (SRS) from SRS Laboratory works by processing existing stereo signals to simulate spatial relationships that a listener distinguishes with the outer ear, adding and subtracting left and right signals, and remixing these signals into the stereo sound. Creative Laboratory Sound Blaster Awe32 sound card, on the other hand, includes a QSound sonic enhancer, which plays audio tracks that have been spatially encoded. It achieves the 3D effect by recombining the left and right channels of information with time delays to create the illusion of sounds in front of or behind the listener.

The success of auditory immersion in VR will depend on the following main factors:

1. The development of realistic computer-generated, non-prerecorded sounds;
2. The development of "sound space" which requires good position and orientation tracking and very complex mathematics;
3. The development of personalised devices that are non-intrusive.

Other factors that should be taken into account in the development of auditory devices in VR are the nature of sound waves and effects; factors influencing sound localisation; and the physiology of hearing.

2.2.3.3. Haptic Displays.

Haptics is the generation of touch and force feedback information. Haptic display devices could be used to make immersion in the Virtual World more believable by recognising that all material objects have properties such as weight, resistance and texture. In addition to this, haptic information could supplement visual and sound information in a Virtual World. For example, haptic texture could be "modelled" for a given object, including such descriptions as heat absorption/reflectivity, bump factor and smoothness/roughness. The success of haptic devices in VR will depend on the following main factors:

1. The development of haptic feedback;
2. The development of computer-generated haptic "realism";
3. The design of non-cumbersome and technically correct haptic devices.

Additional factors which should be taken into account in the development of haptic devices are the physiology of touch; the nature of materials; and the nature of physics.

2.2.4. Communication Devices.

A good VR systems should allow for shared environments. This means two or more people can exist in a VE simultaneously. Not only can these people exist in this environment but they can also interact with each other in different ways. These people could be in the same room or on different sides of the planet. The hardware devices that are necessary to allow people to communicate through VEs are modems, Local Area Networks (LANs), Wide Area Networks (WANs) or even serial cables between two or more machines.

The previous sections have discussed the various hardware components which have converged to form the basis for the different types of VR system available. The next section deals with the different types of software systems which when added to the hardware discussed above form the basis for current VR systems.

2.3. Virtual Reality Software Components.

A variety of products fall under the category of VR software. These range from high-end authoring toolkits, requiring significant programming experience, right down to "hobbyist" packages, for which familiarity with the computer's operating system is the only prerequisite. Despite the differences in the types of VEs these products deliver, the various tools are based on the same VR development principles.

These principles are to create or import 3D objects, to apply behavioural attributes to the objects, and to program the objects to respond to a participant's actions. In this section these various VR software categories will be discussed.

2.3.1. Virtual Reality Authoring Toolkit Software.

VR Authoring Toolkits are the most functional of the available VR software options. Authoring toolkit software can be divided into two categories. The first category is comprised of programs that use programming libraries to provide a set of functions with which a skilled programmer can create VR applications. The second category contains programs with graphical interfaces for creating worlds without resorting to detailed programming. These programs usually include some sort of scripting language in which to describe complex actions. One benefit of the "no-programming-needed" toolkits is that they do allow programming for users who want to customise their applications to meet specific needs. The main advantage of using programming language or scripting languages is that designers can write programs to apply or modify the behaviour of virtual objects. The programming libraries are generally more flexible and have faster renders than the non-programming software, but one must be a skilled programmer to use them. The phrase "professional-level VR" generally refers to the quality of the environments that can be created with authoring toolkits.

The hardware required to run this type of software system varies greatly. Most support a DOS environment with add-in rendering cards, a few work on Silicon Graphics Interface (SGI) and other workstation systems. There are also other packages available that run on vendor specific hardware configurations. The really high-end packages require extremely expensive image generators such as those used in flight simulators. While some toolkits are designed specifically for the PC, most have ported their workstation-based products to this platform.

A standard feature of authoring toolkits is that files may be imported from a wide array of software packages as well as animation scripts and sounds. Most toolkits are capable of editing models to a limited degree once they have been imported. In addition to supporting various software file formats, some toolkits offer graphics-creation capabilities and extensive object libraries.

Some toolkits offer mechanisms for optimising the polygon counts of the objects that will populate the Virtual Worlds. This is extremely useful for designers attempting to design real-time environments as it will help to reduce the number of calculations a computer will have to process and, therefore, speeding up the generation of the VE. This is particularly important when dealing with imported models, which may be too polygonally dense for real-time walkthroughs. Generally toolkits support a level of detail function. Level of detail refers to the process of creating models of various degrees of detail, then programming the simulation to switch among different versions of the same model depending on its visual proximity to the user. This type of function is also a very useful addition when trying to achieve real-time performance. In addition to what the toolkits offer, run-time software can also play a major role in the development of a successful application.

2.3.2. Modelling/Simulation Software.

In addition to VR toolkits, another type of software that enables the creation and display of Virtual Worlds is VR modelling software. VR modelling software provides tools for creating 3D objects and assigning worldly attributes to them. Modelling software varies widely in ease of use and capability. There are some excellent professional packages appearing in this type of software. Most of these systems do not require any specialised hardware beyond the basic computer system. Though less flexible than the toolkits, the lower-end modelling tools offer a number of advantages: they do not require programming, they have a short learning curve, and, compared to the authoring toolkits, are relatively inexpensive. Most offer a rich set of tools for creating 3D objects, or importing them from other packages, and the ability to assign actions, reactions, and connections to them via a point-and-click interface. Also, many programs offer impressive shading and textures, which can go a long way toward enhancing the look and feel of a VE. 3D modelling packages with more sophisticated functionality are proportionally more difficult to learn. Some of the most powerful features of these modellers, such as multiple light sources, mirror surfaces, shading options, and batch animation capabilities, are presently unusable for real-time rendering. Finally, some modelling software packages do not possess adequate flexible object linking mechanisms.

Some modelling packages do not provide the user with refined metrics that allows a designer to create Virtual Worlds scaled to human proportion and this can cause many problems when creating a VE. Some modelling packages do not allow inclusion of sound in the VE which can be a major disadvantage depending on the application that the VE was intended for. However, many of the difficulties associated with modelling software is related to performance - the user must cope with the fact that common operations like displaying the model are relatively slow and laborious. Certain techniques were invented to help, such as multiple viewpoints, but the VE is still inherently being viewed on a 2D screen. You are, however, freed from many of the burdens of the physical world when working with these systems - gravity for example. It is this freedom that allows a designer to model something faster than it can be built.

2.3.3. Walkthrough Software.

At the low-end of the VR spectrum lies walkthrough software. These products allow a participant to move smoothly in any direction while the computer quickly redraws the created world in real-time but many packages do not allow interaction with objects they encounter. There are currently a few fast rendering programs that have been released complete with their source code. These programs are generally copyrighted freeware, which means that the original creators retain the copyright and commercial use is restricted. They are not polished commercial programs, and are often written by students. However, these programs exist to give people a very low cost entry into the world of VR. These programs are the least expensive of all the VR software and are also, by and large, the least effective. In most cases they do not allow the user to import files from, or export files, to other software packages. Nevertheless, such walkthrough software packages are a boon to architects, engineers, set designers, and

Web-page designers. Many computer games that can be considered in this category, from earlier forays such as Wolfenstein up until the latest adventure of Duke Nukem 3D, but these are often closed systems that do not allow much customising or world building by the user.

The main point to be considered when choosing a piece of software is the purpose of the VE and, therefore, the software requirements. Although all three of the VR software categories will benefit from the ongoing hardware developments, the authoring toolkits stand to gain most as they often require high-end expensive image generators.

2.4. The Types of Virtual Reality System.

There are three major types of VR system currently available, Desktop VR, Immersion VR and Third Person VR. The type of VR system chosen has a significant influence on the modes of interaction, the perception of presence, and the experiences obtained in a VE. As a background to understanding VR, and in an attempt to choose the type of VR system which best suits the needs of the tourism industry, it is important to understand the three types of VR that are currently in use.

2.4.1. Desktop Virtual Reality.

Desktop VR, as the name suggests, is based on standard commercially available desktop computers and, thus, uses a conventional computer monitor to display the VE This concept traces its development back through the entire history of computer graphics.

For instance, a Pentium IBM compatible Personal Computer with a powerful graphics card can deliver the speed and reality to make VEs seem real. The user may interact with the VEs with a mouse, keyboard, or any of the controller devices referred to in section 2.2.1.1.

2.4.2. Third Person Virtual Reality.

In Third Person VR the user views himself in a 3D environment. The individual stands in front of a video camera which captures his/her image and body movements and sends this information to the computer which composites the image with computer and laser disc based imagery. The user then views the resulting images - his/her body as seen through the video camera, computer graphics and laser disc imagery - on a monitor or projected on a screen in front of him/her. With this process users are able to participate in a VE by manipulating the computer-generated objects they see in the monitor with their hands, feet and body.

This type of VR system has been evolving since the late 1960s when Myron Krueger, often called "the father of virtual reality," began creating interactive environments in which the user moves without encumbering gear. Krueger's is "come-as-you-are" VR. Krueger's work uses cameras and monitors to project a user's body so it can interact with graphic images, allowing hands to manipulate graphic objects on a screen, whether text or pictures. This is the projected image display referred to in section 2.2.3.1. The interaction of computer and human takes place without covering the body with equipment and, therefore, the burden of input rests with the computer. Cameras follow the user's body, and computers synthesize the user's movements within the VE. The computer constantly updates the interaction of user's body and the synthetic world that the user sees, hears, and touches. In Krueger's Videoplace, people in separate rooms relate interactively by mutual body painting, free-fall gymnastics, and tickling. Krueger's Glowflow, a light-and-sound room, responds to people's movements by lighting phosphorescent tubes and issuing synthetic sounds. Another environment, Psychic Space, allows participants to explore an interactive maze in which each footstep corresponds to a musical tone, all produced with live video images that can be moved, scaled, and rotated without regard to the usual laws of cause and effect.

At least one commercial system uses this approach, the Mandala system illustrated in Figure 2.10. This system is based on a Commodore Amiga with some added hardware and software. A version of the Mandala system is used by the satellite and cable TV channel Nickelodeon for a game show, Nick Arcade, to put the contestants into what appears to be a large video game.

2.4.3. Immersion Virtual Reality.

Immersion VR describes a system that "immerses" or surrounds the participant in an environment. The participant can hear, see and, perhaps, feel nothing other than the artificially created environment In order for an individual to participate in an immersion VR system certain equipment is deemed necessary. Some of the equipment which one would expect to find as part of such a system would be HMDs, bodysuits and gloves which track and project a participant's movements while also allowing them to manipulate objects within the VE. As a result the participant actually perceives that they are inside an environment. Immersion VR systems are the most demanding type of VR system in terms of the technology required to offer the appropriate system response and graphics display rates.

According to this view, VR means sensory immersion in a VE. Such systems were first popularised by VPL Incorporated. The HMD cuts off visual and audio sensations from the surrounding world and replaces them with computer-generated sensations. The body moves through artificial space using feedback gloves, foot treadmills, bicycle grips, or joysticks. A prime example of immersion comes from the U.S. Air Force, which first developed some of this hardware for flight simulation. The computer generates much of the same sensory input that a jet pilot would experience in an actual cockpit. The pilot responds to the sensations by, for instance, turning a control knob, which in turn feeds into the computer, which again adjusts the sensations. In this way, a pilot can get practice or training without leaving the ground. To date, commercial pilots can upgrade their licenses on certain levels by putting in a certain number of hours on a flight simulator. Computer feedback may do more than re-adjust the participant's sensations to give a pseudo-experience of flying. The feedback may also connect to an actual aircraft, so that when the pilot turns a knob, a real aircraft motor turns over or a real weapon fires. The pilot in this case feels immersed and fully present in a Virtual World, which in turn connects to the real world. When one flies low in an F-16 Falcon aircraft at supersonic speeds over a mountainous terrain, the less one sees of the real world, the more control one has over the aircraft. A virtual cockpit filters the real scene and represents a more readable world. In this sense, VR can preserve the human significance of an overwhelming rush of split-second data. The heads-up display in the cockpit sometimes permits the pilot to view the real landscape behind the virtual images. In such cases, the simulation is an augmented rather than a virtual reality.

A nice variation of the immersive system uses multiple large projection displays to create a 'Cave' or room in which the viewer(s) stand. An early implementation was called "The Closet Cathedral" for the ability to create the impression of an immense environment within a small physical space. The Holodeck used in the television series "Star Trek: The Next Generation" is a very long-term extrapolation of this technology.

2.5. Which Type of Virtual Reality System is the Most Appropriate to the Tourism Industry?

VR has an increasingly important role to play in the tourism industry, offering real-time visualisation and interaction within Virtual Worlds. In this scenario, however, Immersion VR and Third Person VR systems currently have a number of drawbacks. Participants can suffer from what is commonly called simulator sickness while using headsets. Simulator sickness is a disturbance produced by simulators, ranging in degree from a feeling of unpleasantness, disorientation, and headaches to nausea and vomiting. Many factors may be involved, including sensory distortions such as abnormal movement of arms and heads because of the weight of equipment; long delays or lags in feedback, and missing visual cues from convergence and accommodation. In both these systems frequent recalibration of the equipment may also be necessary.

Desktop VR is not significantly affected by any of these problems. Furthermore, the fact that interaction with VEs is controlled from the "desktop" immediately makes Desktop VR suitable for multiple viewing at presentations. The absence of a headset also overcomes the problem of the clarity of the screen, giving a greater level of resolution and detail. Using industrial-standard computer hardware also ensures higher reliability, lower maintenance costs and, most importantly from a tourism perspective, a larger diffusability through the industry. Though the resolution and accuracy of headsets are bound to improve, they do not reflect the needs of the serious user. These advantages of Desktop VR, both for VE authoring and visualisation, make it the most suitable solution for the tourism industry. Consequently, the specific application described in this dissertation will be based on a Desktop VR system.

2.6. Summary.

In this chapter the types of VR systems currently available are examined in an attempt to choose the most appropriate type for the specific application of VR to the tourism industry designed in this research. There are at present three broad types of VR system; Third Person VR, Desktop VR and Immersion VR. In order to choose which type is the most appropriate to this specific application or, indeed for any VR application, one must first examine the specific hardware and software components which come together to form the three different types of VR systems.

The specific hardware components that form the basis for VR systems at present may be divided into four categories; transducers, image generators, display devices and communication devices. Transducers are input devices that communicate a participant's actions to the VE. These actions may be movements, speech or brain activity. Movement transducers are the most common type of input device and, therefore, several types of movement transducers are discussed in this chapter. Image generators are the processors used to generate the VE. There are three types of image generators - visual image generators, audio image generators and haptic image generators. All three types of image generators and the importance of each type were discussed. Display devices are the mechanism through which a participant views what occurs in a VE. There are also three types of display devices. These are visual displays, audio displays and haptic displays. As sight is our most important sense the main focus of this section was placed on visual displays. Finally, communication devices are discussed. These are devices that allow two or more participants to exist in a VE simultaneously. These participants could be in the same room or on other sides of the world. These devices vary from simple serial cables to modems to networks.

VR software systems fall into three categories; VR authoring toolkits, modelling/visualisation software and walkthrough software. VR authoring toolkits, also referred to as "professional-level VR" are programs which usually include some sort of scripting or programming language which enables a designer to construct a highly vivid interactive VE. The hardware required to use these toolkits varies greatly from commercially available PCs to extremely expensive image generators such as those used in flight simulators. Modelling/visualisation software varies widely in ease of use and capability. Most modelling packages do not require any additional computer hardware beyond a basic PC. This software type does not normally require programming, normally has a shorter learning curve than toolkits and, compared to the authoring toolkits, is relatively inexpensive. The final software category, walkthrough software, includes packages found at the low-end of the VR spectrum.

Desktop VR is considered to be the most appropriate system in the context of this research. Desktop VR is selected because it is not affected by the problems associated with both Third Person VR and Immersion VR. The fact that Desktop VR is based on commercially available PCs also ensures higher reliability, lower maintenance costs and a larger diffusability through the industry.

Chapter Three

Applications of Virtual Reality -

"From the Classroom to the Boardroom"

3.1. Introduction.

The VR Industry is at present where the PC industry was perhaps 20 years ago. After the introduction of the PC, applications gradually began to emerge in business, industry, education and research. VR is still in a nascent, "learning" stage, yet, technology has advanced to the point where it can be used as a practical tool in select applications. The potential use of VR in a wide range of applications is driving VR development despite the high implementation costs. VR technology is currently diffusing into a number of application areas such as information manipulation in medicine, engineering, architecture, rehabilitation, scientific visualisation, telecommunications and entertainment (Rheingold 1991). This diffusion is being promoted by the recent improvements in technology combined with a downward trend in hardware costs. It is reasonable to expect that the technology may diffuse more quickly in applications where 3D rendering and spatial cognition are required features. Most applications development is still heavily concentrated in University and government laboratories, especially in relation to medicine and aerospace, rather than in commercial enterprises. Work is consequently focused on developing pre-production prototypes and demos. In this chapter the areas in which VR is currently being employed are discussed but in an attempt to place these in perspective, the main tasks involved in VR applications development are first outlined.

3.2. The Main Aspects Involved in Virtual Reality Application Development.

The potential of VR lies in the prospect of allowing a participant the freedom to explore Virtual Environments (VEs). These VEs may differ greatly depending on the tasks that the participant is allowed to perform. The tasks performed by users vary according to the characteristics of the application domain. When considering VR, three domains are identified, Figure 3.1. They are:

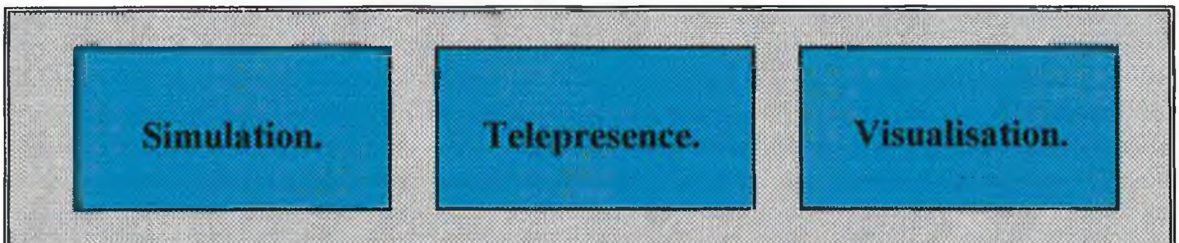

Figure 3.1. The Three Aspects Involved in Virtual Reality Applications Development.

Simulation: Representing real-world environments, perhaps with a degree of fidelity that is almost indistinguishable from the real thing. This is currently used in automotive construction and flight training simulators.

Telepresence: Representational user inclusion in image space in a form that can realistically mimic and reflect user movements, desirably with realistic sensory feedback. Typically applications are to be found in telematic robotics in hazardous environments.

Visualisation: Enabling the user to interact with abstract, conceptual or "customised" environments created from data that are not normally directly accessible to the human senses such as molecules, generic information, other worlds and games. Visualisation is also extremely important in areas such as engineering and architectural design where schematic representations, drafting aids and 3D displays and models have developed over many decades. A user may also take on the bodily appearance, sensation, and capabilities of another object or being as in role-playing adventure games.

In simulation, the style of interaction is intended to map directly onto that adopted in the "real" world. An example of such a scenario would be what a pilot sees outside and interacts with inside the cockpit: object representations and user interactions are exactly the same as the real thing. In the case of telepresence, there can be deviations from direct mapping. For example, a robot's cameras may provide a much lower viewpoint height than the user's eye. In addition, novel user tasks can be created that have no real-world equivalent such as walking through the core of a working nuclear reactor. Finally, in the case of visualisation, there may be no corresponding real-world task at all, requiring the design from scratch of tasks, interaction styles and display objects to best enable user performance. What makes this interesting is the underlying philosophy that user interaction can be designed from the start to feel natural and intuitive.

3.3. The Applications of Virtual Reality.

The now popular concept of VR as a household item offering unlimited entertainment in "designer fantasies" is unlikely to reach the mass market for some time. The more immediate and interesting applications of VR are practical ones in the fields of science, medicine, industry and technology, in which the professional user may benefit from an enhanced experience of working in a new dimension in a more-or-less familiar environment. Some examples are considered.

3.3.1. Architecture.

Computerised drafting aids and visualisation techniques have been used by architects for the past twenty years. Over the years, professionals in the fields of design engineering and architecture have been quick to appreciate the advantages of 3D technology for savings in development costs and time. The computer was used as a means to externalise the design ideas. Architects who became skilled with these tools felt empowered to challenge even their own creations and were encouraged to try out new ideas and more actively include their clients in the planning and design process.

However, even the best of the current packages leaves the user an outsider, denying him the experience of interaction with the environment. Also, due to the design portrait being static, changes to the design are costly and time consuming, often requiring, regeneration of several frames to create the images or animation. VR is the logical next step in the development of such tools. VR enables architects to both move around and interact with the environment, and to manipulate it in real-time. This is an ideal display environment to help clients to understand precisely what buildings will look like once they are constructed and, in this way, to include them in the design process.

Most virtual architectural applications rely almost solely on visual exploration and require, little if any, output to any of the other senses and are therefore referred to as "walkthroughs". Walkthroughs facilitate the collaboration between the Designer-Client-Contractor loop. In order to thoroughly discuss the application of VR to architecture it is useful to concentrate on three distinct aspects of VR: Virtual Past, Virtual Analysis and Virtual Design.

Virtual Past: Representations and reconstructions of past architecture, either as replicas in the form of drawings, watercolours or computer models, are well known. They form an important part of design education. Whereas in most cases the form of nonexistent or destroyed architecture can be restored, other known information and characteristics of the period are often unavailable and neglected. Making this information available in its most appropriate form is possible in a VR environment, with the inclusion of acoustic, colour and atmospheric impressions. All historical reconstructions are impaired by the fact that time dependent attributes such as scales, political opinions and societal values to the artifacts, buildings or sites under construction are applied. Stenvert (1992) has attempted to study the effects of constructing the past with the aid of computers, which serves as an interesting approach which can be adapted to VR. In order to construct a more realistic reconstruction, researchers should be able to step back into time with the aid of a Virtual Environment (VE), which must closely resemble conditions at the time the object was constructed. Most of these conditions can be simulated directly as 3D objects or made available upon request as multimedia information. The goal is to equip the researcher with as much knowledge of the past as possible in order to reconstruct it.

Virtual Analysis: Design analysis is more difficult than generation. A comprehensive analysis of an existing building almost requires a designer to retrace the original design process to understand the design influences. VR facilitates such analysis, in that it lets the designer experience simultaneous views of a building and related analysis data. Analysis of only one attribute of a building can be performed efficiently by computers. A major source of problems in construction is the misinterpretation of the 2D working drawings. The projection of a true scale "skeleton" of a building and all of its components could resolve most 3D questions.

In summary, virtual analysis is an appropriate tool for testing design models which are inherently geometric, acoustic, energy, or based on a quantitative or computable nature. VR is a tool for conceptualising otherwise invisible properties of a design.

Virtual Design: Virtual design is the application that interests architects most (Schmitt 1993). It requires the definition of a new working environment for architects. VR provides the physical and technical means to continue a well-developed direction in an architectural theory - futuristic scenarios.

Virtual design can be simulated in a VE using more design-influencing information than is available today. The early design phase and the preparation of a building program may serve as an example. In this process one of the first steps is a site visit. While many constraints will restrict this important task to a minimum amount of time, tele-sensoring and tele-surveying could be of help (Balaguer & Mangili 1991). This involves the placement of an instrument on the site which will stay there during the design process. The instrument has the following two purposes:

Data collection. Temperature, atmospheric and lighting conditions, precipitation, wind and insulation are monitored and compared with average data of the area. Differences are detected and used to define a more accurate microclimate of the site. A design workstation receives the data, which serves as a realistic input to the design process.

Visual control. The instrument is equipped with a stereo camera that allows the designer to request particular views from the site. The views can be used to create simulations of the new building and the existing environment.

Ideally, the instrument is movable, teleoperated from the designer's workstation. In order to conduct a constant evaluation of the new design in the given context, selected views of the site during different times on the day and year are necessary.

The three areas of virtual analysis, virtual past and virtual design are the main architectural areas VR applications in the foreseeable future. These areas cover the three aspects of using VR; assisting the architect in designing a building, evaluating the building when it is being built and reconstructing an old building in order to learn from the design.

3.3.2. Medicine.

VR technologies are believed to possess the potential to revolutionise not only future physicians and other health professionals, but also to keep practitioners abreast of developing information and procedures, provide a less subjective, common basis for evaluation, as well as to actually facilitate the execution of certain procedures. The number of potential VR applications within the sphere of medicine alone is seemingly boundless. However, the most obvious application for the new technology are surgical simulators used for training and as preparation for surgeons about to perform serious operations. Therefore, the application of VR as a surgical simulator will be discussed first.

Surgical Simulators: A significant factor in medical education is the high costs associated with the training of surgeons. Studies have repeatedly shown that an alarmingly high number of mistakes are made by surgical residents during their first surgical undertakings on live patients (Merril et al. 1995). Referred to as the "learning curve", many medical specialists now believe that such mistakes may be averted through the use of VR, in the guise of surgical simulators.

Surgical simulation, like flight simulation, allows the user to be trained to perform a complex task using an interactive computer environment. This interactive environment has progressed from a 2D screen to 3D VR. Current research endeavours in medicine involving the use of VR technology has concentrated on three critical aspects of surgery, including the provision of greater, more precise control over surgical instruments, allowing surgeons to better visualise the surgical site, and endowing the surgeon with the ability to simulate procedures prior to the actual operation (Schratt et al. 1994).

A surgical simulator has three basic components: the computer, the interface and the physical model. The computer and the types of interface used are quite similar to those used in any VR system and are discussed in some detail in Chapter 2. It is in the area of the physical model that surgical simulators differ from other types of VR system. A detailed physical model of the entire body is needed to develop a realistic surgical simulator. Specific parts of the body that include the skeleton with attached tendons have been developed. One of the first was a model of the hand that allows realistic tendon transfer for patient with nerve palsies. A newly developed model of the lower extremities, which allows predictions to be made about the outcome of reconstructive surgery, was first used for osteomies around the hip joint and has progressed to muscle/tendon reconstruction. This model was used to demonstrate the usefulness of performing surgery within a 3D VE instead of on a 2D screen. The demonstration provides a test-bed to study the different approaches to visualisation and interacting within a VE.

Surgical simulators allow surgeons not only to practice a certain procedure repeatedly prior to actually entering an operating room, thereby serving to greatly reduce the risk to the patient, but also to program in to the operation any number of potential complications, serving to better prepare the surgeon should such an emergency arise during surgery. In addition, a large amount of a surgical students' time is wasted waiting to observe a certain surgical procedure, especially those which may be infrequently performed. Due to advances in VR technology an operation can be performed on a surgical simulator at any time, thus, saving valuable time. Finally, it is believed that the use of VR would serve to greatly bolster the confidence, and accelerate the learning curve, of the student (Larijani 1994).

Computer-based surgical simulators, like VR, are still in their infancy. Each of their components - the computer, the interface and the physical model will have to undergo major advances over the next few years in order for the simulator to receive overall acceptance. The computer graphics engine will support a more realistic VE, while the HMDs and the whole-hand interface device will enable the surgeons to see and feel the more realistic environment. However, the physical model of the patient must undergo the most remarkable advances. The integration of the components of the model will result in a computational representation of the human body that includes all of its key parts: bone, muscle and skin. In addition, the functionality of these parts will be physically modelled in more and more detail, moving static to dynamic models.

Disability: A great deal of research is currently being directed towards the use of VR to aid those who are afflicted with Parkinson's disease. Although the administration of hormone therapy has been the long-standing course of treatment, a side-effect of such drug therapy is erratic and uncontrollable body movements. Recent research into the disease has revealed that some of the people suffering from Parkinson's are able to respond to visual cues placed on the floor in front of them in such a way that they are enabled with the ability to step over these objects. If objects are placed sequentially in front of these patients, they are essentially able to walk over them. Although much research has gone into investigating the physiological cause of this phenomena, few answers have been found. Researchers now believe that VR technologies, namely, the HMD, can be used to continually superimpose objects on the floor in front of the Parkinson's patient, thereby empowering them with the ability to walk (Bennett 1995).

VR is also being demonstrated as a powerful means of aiding others who are physically disabled. VR systems have also been demonstrated to be an invaluable teaching tool for those attempting to learn how to navigate about their world, such as those confined to a wheelchair (Pimentel & Teixeira 1992).

Rehabilitation: Research is also being conducted on new rehabilitation devices for patients. "GloveTalker", a data glove sign-language device, allows the user to speak, with the use of a DataGlove, by signalling the computer with their personalised set of gestures: the DataGlove is used to recognise the user's hand positions and this information is passed through to the computer's voice synthesis system, which will speak for the DataGlove wearer. In effect, GloveTalker interprets signs and then translates them into commands that appear on the computer screen as written language, as well as synthesized speech. The voice output can be sent easily over a computer network or over a telephone system, thus enabling vocally impaired individuals to communicate verbally over a distance. The system ability to translate crude motor functions into computer synthesised speech, has opened up the ability of communicatively impaired individuals to interact with others.

Psychiatry: Preliminary studies in the area of psychiatry indicate that VR is an effective medium for the treatment of a variety of disorders, including arachnophobia (fear of spiders) and acrophobia (fear of heights). Using a VR system, the patient is able to control both the intensity and duration of exposure to stimuli (either spiders or heights). Researchers have found that self-reported anxiety, fear, and avoidance were reduced significantly in subjects exposed to VR graded exposure treatments. Thus, they concluded that VR was an effective means by which such fears could be reduced. As such, a growing body of professionals believe that VR may offer a time- and cost-effective way to conduct effective exposure therapy in the treatment of a variety of such disorders (Rothbaum et al. 1995).

Radiation: One of the most difficult aspects of intensive cancer treatment for many physicians is involved in radiation treatment planning - positioning powerful beams of radiation so that they impinge upon the malignant tumour without damaging the surrounding healthy tissue. Traditionally, diagnostic radiologists and radiation oncologists have attempted to plot the 3D trajectories of these beams based solely on 2D data obtained from X-ray images. However, by using VR technology, doctors are now endowed with the capability of using 3D patient data obtained from computerised axial tomography (CAT) scans and magnetic resonance imaging (MRI) scans to run trials of a variety of beam placements in order to determine which course of action will yield the best possible effect. Using a 3D image of the patient's body, the physician can set a path for the beam and examine it from a variety of angles. Low-frequency tones could be used to provide feedback as to how far the beam is off-target from the tumour, while colour changes displayed by the system could be used to indicate when healthy tissue was being impacted (Pimentel & Teixeira 1992).

Reconstructive surgery is another area that VR technology promises to revolutionise. Large wounds, for example, often require surgical reconstruction that necessitates transferring healthy muscle from one area of the body so that it may be used to cover the affected area. Generally, surgeons have only their experience to rely upon in choosing which muscle to relocate so as to minimise the impact of donor tissue removed on other parts of the body. Such a procedure involves a series of trade-offs between the various groups of muscles - moving a muscle to the affected area will necessarily result in a reduction in strength from the site of procurement. By using a VR system, surgeons could use actual patient data to augment their understanding of the bio-mechanical consequences of choosing various muscles so as to get the best results from the surgical procedure. It is believed that such an application will aid physicians in all phases of reconstructive surgery, including preoperative planning, training, and, ultimately, surgical assistance during the actual operation (Pimentel & Teixeira 1992).

Dentistry: To date, the field of dentistry represents a realm in which individual VR components are being more readily utilised than the total experience provided by a complete system. VR appears to be largely confined to relaxation therapy devices which involves the use of HMDs equipped with stereo earphones to alleviate patient anxiety (Digital Image FX 1995). These HMDs are connected to videocassette recorders and produce images, generated from the patient's choice of movie, television show, or computer game, which appear to float in front of the patient. Proponents of this equipment claim that it is an effective means by which to alleviate patients' fear and anxiety by removing them psychologically from the dental procedure being performed, and replacing it with a pleasant experience (Swarden 1994). The accompanying sound provided by the stereo earphones serves to mask the sounds of instruments being used during the procedure.

Many dentists who have introduced this technology into their practices claim that because the headset allows the patient to relax, they are able to perform procedures more quickly, thereby increasing the number of patients that they are able to see on a daily basis (Swarden 1994). In the near future, some dentists plan to begin using the headsets for educational purposes by demonstrating procedures and proper dental hygiene to their patients. Optional cameras may also be purchased for the device which would allow the dentist to temporarily interrupt the program being viewed in order to show the patient exactly what was going on inside their mouth, if they so desired (Andreassis 1994). Orthodontic reconstruction would appear to be an obvious procedural future application for VR.

Health Studies: To date, the literature provides little evidence of VR applications that have been used in other health professions including physiotherapy, occupational therapy and nursing. However, it seems likely that VR applications currently employed to create walkthroughs to facilitate architectural design would also offer decided benefits to occupational therapists who are frequently faced with the task of how to modify a patient's home or workplace following what is usually a reduction in mobility so as to offer them the greatest freedom and independence. Using such a system, the therapist could model the individual's environment, make the necessary modifications and allow the patient to interact with that environment with the aid of immersion VR in order to evaluate the effectiveness of the modifications (Pimentel & Teixeira 1992).

In addition, preliminary studies indicate that games employing the use of a DataGlove that allows the player to manipulate imaginary, weightless objects, may be used to augment or replace conventional physiotherapy exercises. Researchers have found that when such VR applications are used, patients worked harder at their exercises thereby serving to facilitate their therapy (Mestel 1993). VR applications may also be used to track the progress of patients undergoing rehabilitation therapy. These patients generally progress through a series of small improvements in motor skill acquisition and function. Often, these improvements are too small for the patient to realise that they are, in fact, making progress, and thus, they often become frustrated and disinterested in the course of their therapy. VR systems are able to track even the slightest improvements and may thereby serve to bolster the patient's confidence and motivation (Pimentel & Teixeira 1992).

The potential also exists to employ immersion VR technology to conduct ergonomic analyses of working environments in an attempt to pinpoint patterns of behaviour that may lead to or exacerbate stress or strain injuries such as Repetitive Strain Injury (RSI) (Larijani 1994). Such conditions are commonly observed in workers who perform tasks that require limited and repetitive hand movements. RSI is currently considered to be one of the most prevalent and debilitating work-related injuries (Pimentel & Teixeira 1992). However, VR now offers the potential to allow occupational health care researchers to discover ways in which the frequency and occurrence of such injuries can be minimised.

The ability of VR to similarly benefit other health professions has yet to be explored and remains an area where great potential is yet to be fully realised. Areas currently being explored for the potential application of VR technology include stroke and injury rehabilitation, occupational assessment, and home therapy (Greenleaf 1996)

The potential of VR to significantly alter the nature and enhance the quality of medical education remains a much-discussed topic at the professional meetings of various medical specialties. But concern is still expressed regarding the technical obstacles blocking the technology's acceptance in this arena, such as less-than-realistic imagery and the absence of convincing haptic simulation.

3.3.3. Education/Training.

Computer graphics technology enables us to create a remarkable variety of digital images and displays that, given the right conditions, effectively enrich education (Bricken 1990) Real-time computer graphics are an essential component of the multi-sensory environment of VR. This section addresses the unique characteristics of emerging VR technology and the potential of VEs as learning environments.

Although military, government, and industry training programs have employed simulations using some form of VR for years, the development of VR applications suited to educational uses is largely unexplored territory.

The case has been made that psychological processes that become active in VR are very similar to the psychological processes that operate when people learn through interaction with objects and events in the real world. Put simply, VR promises to be an extremely useful experiential learning tool in education. The idea is simple, i.e. everything one does to educate with words and with pictures can be provided as a virtual experience. Theoretically, VR may be used in the context of many subjects and appeal to different learning styles. The traditional method of text based instruction may be augmented by the experiential method of VR. Furthermore, the classroom use of VR could be varied. On one end of a continuum is the learning that occurs from building a world. For students, it is a synthetic application of knowledge about the subject, it allows them to reflect on information gathered as they explore the world in VR, while also being an empowering creative process. On the other end of the continuum, the students enter a pre-fabricated world where they learn from the exploration and not the creation of the world. Of course, any compromise between these two extremes is possible. Perhaps a student could, for example, build part of a world or add on to an existing world in order to facilitate learning.

Current VR systems provide new capabilities for perceptual expansion, for creative construction and for unique social interactivity (Bricken 1991). These characteristics of VR are relevant in five areas of educational theory:

1. Intuitive Characteristic,
2. Experiential Characteristic,
3. Constructive Characteristic,
4. Social Characteristic,
5. Customising Characteristic.

VR offers teachers and students unique experiences that are consistent with successful instructional strategies: hands-on learning, group projects and discussions, field trips, simulations, and concept visualisation. Within the limits of system functionality, it is possible to create anything imaginable and then become part of it. The VR learning environment is experiential and intuitive; it is a shared information context that offers unique interactivity and can be configured for individual learning and performance styles. These are the characteristics of VR that make it such a suitable tool for education and, thus, deserve further discussion.

Intuitive Characteristic: At the intuitive level, VR mimics they way that humans have learned to interact with the real world. A VE empowers a participant to move, talk, gesture, and manipulate objects and systems in a natural way. The skills needed to function within a VE are the same skills practised in the physical world. This method of representing and interacting with information is fundamentally different from the way computers have been used up until now. Novices require minimum accommodation time (Bricken 1990). Skilled users can represent and manipulate increasingly complex information in forms that are easy to remember and interpret (Furnass 1988).

The intuitive nature of VR means that the computational environment removes symbolism in favour of displaying information in an innately recognisable form. Rather than teaching a structure of symbols, such as algebra, and then teaching the meaning of that structure, in VR meaning is taught through experience and then, if necessary, symbolic links are made. Manuals and written descriptions may be integrated into the simulated display of objects. More fundamentally, written materials may be unnecessary when replaced by direct experience with virtual objects. Therefore, using VR as a learning vehicle may drastically reduce the requirement of learning abstract concepts, such as written language or jargon of a particular field, in order to convey the message being taught.

Experiential Characteristic: VR is an experiential medium and it is this experiential quality of VR that is fundamental to the learning process (Papert 1980). A VE is a place where participants can have any number of different learning experiences. Knowledge construction arises from first-person experiences that can never be entirely shared. VR allows first-person experiences by removing the interface that acts as a boundary between the participant and the computer. In this, VR technology is unique. It alone allows a synthetic experience to capture the essence of what it really means for a person to come to know the world. Participating in a VE allows us to construct knowledge from direct experience, not from descriptions of experience. Any learning that is mediated by a symbol system, whether text or spoken language, is inevitably a reflection of someone else's experience not our own. Any requirement that one uses a symbol system to communicate about the world we have constructed to someone else can never permit that other person to know our world as we know it.

Constructive Characteristic: The word "constructivism" is used to describe learning that arises from physical interaction with objects in the real world (Papert 1991). With some exceptions (Bricken 1991;Bricken & Byrne 1993; Bricken 1990; Winn & Bricken, 1992) educators have not made the connection between constructivist theories of learning and VR, thereby missing the opportunity to provide a theoretical basis for applying VR in education. The characteristics of VR and the axioms of constructivist learning theory are entirely compatible and constructivist theory provides a valid and reliable basis for a theory of learning in VEs. VEs are constructive environments in which participants can create, manipulate and edit any form of digital information. Constructivists, rather than prescribing learning outcomes, focus on tools and environments for helping learners interpret the multiple perspectives of the world in creating their own (Bricken 1990).

Social Characteristic: Human learning presupposes a specific social nature and social process (Vygotsky 1978). Education is inherently social. VEs can be considered in both individual and social contexts. VEs exist in a social context in that they can be networked to provide shared environments that allow communication and collaboration between local or distant participants. The ability to literally exchange or share points of view in multiple-participant VEs may intensify this social learning experience (Winn 1993). Co-creating VEs for learning allows teachers and students to use computers effectively in a work-group situation.

Research in collaborative learning abounds with evidence of its educational value: The learner tends to be more productive in a group situation than working in isolation. Ongoing discussion involves the active participation of students in the teaching-learning process; attention is then directed toward the learning activity (Bricken 1990). One of the main reasons for using shared VEs in education is to allow students to solve practical tasks with the help of their speech, as well as with their eyes and hands.

Customising Characteristic: VR provides the potential for completely customised, individualised learning. Educational environments will uniquely respond to the participant, in terms of both needs and preferences. A student model will not be necessary, instead the teacher and student will modify the environment in support of student behaviour. Teachers may in the future represent information in forms that are most compatible with a student's particular learning style, selecting interactivity options that match student performance characteristics. Tools for movement and manipulation within the VE may be configured to the physical needs of the individual and the requirements of the task.

The practical potential of VR is still being explored. Of the number of application areas that suggest themselves, education is clearly worth immediate investigation. VR was devised to enable people to deal with information more easily, and it has been successfully developed to facilitate learning and task performance for over 20 years in the U.S. Air Force (Furness 1978). Public education and training applications are a natural extension of this work.

Among those commercial companies pioneering the effort are Motorola and Northern Telecom, which have traditionally invested heavily in classrooms and training centers, and have recently begun to migrate their training efforts from traditional methods to desktop VR software that simulates hands-on work experience. Motorola and Northern Telecom found that their low-end VR training software is not only cheaper and easier to deploy for geographically dispersed students but, in general, also yields better results. Motorola tested VR in 1994 before implementing a pilot program. The test compared a sample of people trained with VR with those trained in the classroom. The results showed VR students outperformed the others once they were placed on the job. Another advantage is the capability to portray working scenarios that may not currently exist or to simulate critical experiences without real-life risks. What-if scenarios, positive and negative reinforcement regimes, records of behaviour and of common errors, weaknesses in the training sequence, lessons learned, in fact all aspects of training evaluation can be automated and recorded.

Any new technology obviously needs thoughtful introduction into classrooms. Technology does not, by itself, improve education, and even the most promising educational innovation needs skilful application to be effective. The use of VR as a learning environment requires an understanding of the shift from education as primarily a text-based experience to one that is multi-sensory. VR, from an educational perspective, is less formal than traditional classrooms, perhaps more fun, and more realistic in terms of the wealth of experiences that might be generated. The environment, when effective, invites user participation in problem solving, concept development, and creative expression. Though we can only speculate on the contributions VR will make to education, it seems, from emerging evidence that students will participate in responsive environments in which they will become engaged in full body-mind learning (Ferrington & Loge 1992).

VR has a definite and specific role to play in education. In any given potential application however, the appropriateness and relative cost-benefit of introducing this new, and currently expensive, technology should be carefully considered. Essentially, VR is a hyper-medium best used in an individual or self-learning environment and, as such, complements less expensive, programmatic alternatives such as film.

VEs extend the horizons of the learning arena beyond the classroom walls, providing students and teachers with yet another, set of mind tools. Training modules can exploit the mentality of active minds and the heightened interest levels indicative of VEs to provide ideal venues for discovery and self-learning. The ability to interact with the environment and its displays allows a person to try courses of action not possible in the real surrounding. The increased number of "what if" choices available in VR encourages curiosity and stimulates the flight of thought so necessary for creative work. Discriminate selection of applications for this use will provoke new ways of thinking and act as a catalyst for learning opportunities.

3.3.4. Manufacturing.

"Virtual Design" is the use of VR systems for the computer aided design of components and processes for manufacture. VR is used extensively in engineering design, in the production of virtual prototypes, and in creating a cost-effective and efficient bridge between development and production function in industry. It offers unrivaled scope for creating and viewing 3D engineering models, later to be passed to numerically-controlled machines for real manufacturing. Practical and efficient use of Virtual Manufacturing Technology is a necessary step as more and more emphasis is placed on zero defect manufacturing.

Major companies are maintaining their global competitiveness by designing products better and faster with VR. Boeing, for instance, developed a VR model of their 777 aircraft for this purpose and General Motors is using a virtual prototype rather than a physical model of one of its new model cars to evaluate the interior design for aesthetics, engineering, safety, and ergonomic features.

VR has been applied to a wide range of problems associated with industrial maintenance and manufacturing. Most applications can be placed in one or more of six main categories: visualisation of complex data, controlling industrial robots, remote operation of equipment, enhancing communications, operations training, and virtual prototyping and design. These applications are but an initial step in identifying opportunities for using VR to improve the competitiveness of manufacturing organisations.

In particular, VR technologies can play a substantial role in the interactive design and analysis of complex production systems and processes. The 3D nature of VR, and the mechanisms for interacting with objects in VEs, makes the enormous amount of technical data available much more accessible to decision makers; travelling through, and manipulating objects within, the virtual facility offers a much more natural and direct interaction than working indirectly through programs providing only 2D representations of the problem.

VR has excellent effect in transforming flat architectural drawings into realistic 3D designs. Designers or engineers eventually study, test and make changes on the virtual prototype. Using VR can shorten the new product development cycle and reduce the cost of production since there is no need to go through rounds of costly prototyping.

3.3.5. Entertainment.

Early developments of VR were closely linked to the Film and Entertainment Industries, especially the computer and video game markets. Film distributors and producers have experimented with the various imaging technologies in the quest to offer the complete experience for consumers of entertainment. Advances such as 3D glasses, sound and high quality imaging technology within cinemas have emerged and production teams have recently created some very realistic computer-generated graphics and special effects in movies such as "*Lawnmower Man*", "*Johnny Mnemonic*", and "*Jurassic Park*".

Entertainment was the first application for which VR systems were manufactured in quantity. These systems are often used as single units in an arcade environment or networked together in groups for team experiences or competitive events. The first of these experiences was based on flight simulations in which a player used console controls to fly and shoot down the opposition in order to score points. The appearance of VR on the leisure market caused great public interest and attracted the attention of the media. The resulting spate of documentaries and news programmes spread the interest in the technology to a much wider audience and raised the public profile of VR all over the world.

At present, VR is used primarily as an entertainment device. The most successful companies and products are directed towards this end, but this does not mean that these are trivial systems, on the contrary, they typically make use of a wide range of VR components, including fast rendering engines, HMDs, multi-user worlds, and high levels of interactivity. One of the major limiting factors facing VR at present is the high cost of the technology. Increasing volume output will produce a reduction in price and many more organisations will soon be able to afford VR. Therefore, the lessons learned in developing these types of systems will help propel the entire industry.

3.4. Summary.

The technical ability to represent environments electronically via VR is creating a great deal of excitement in a variety of academic, governmental, and commercial settings. So far, the primary application areas are in entertainment, architecture, medicine, education/training and manufacturing.

The specific applications range from flight simulators for NASA's space shuttle program, to arcade games, to architectural renderings in CAD, to medical image projections for planning radiology for cancer. The scope of VR will broaden and many other application areas will be effected by its power. Among these other areas that could benefit enormously from the application of VR is the tourism industry. It is quite surprising that, currently, there is little emphasis placed on the use of VR in the tourism industry. The following chapter addresses the issues associated with the introduction of VR to the tourism industry and will uncover the possible areas where VR applications are most likely to occur in the tourism industry in the future.

Chapter Four

An Application of Virtual Reality to the Tourism Industry

4.1. Introduction.

In this dissertation so far the issues of VR, its component parts and its industrial applications have already been addressed. In this chapter the possible applications of VR to the tourism industry are discussed and the design considerations for a specific application are laid out. In order to develop an application for any industry a clear understanding of that industry must be obtained. Once this is achieved the next step is to develop an application to suit the needs of the Industry. The Industry in question here is the tourism industry and this chapter begins by discussing the tourism industry and, more specifically, the relationship between tourism and technology. This is absolutely essential in order to understand how the use of VR may impact on the tourism industry in the future. The second part of this chapter explores the design considerations of a specific VR application to the tourism industry.

4.2. The Tourism Industry.

The tourism industry is comprised of many sectors, groups and organisations, and to date there is no internationally accepted definition. However, the United Nations Conference on International Travel and Tourism defined tourism as:

> *the consumption, production and distribution of services for travelers who dwell in some place other than their domiciles or workplace for at least twenty four hours. Shorter sojourns are regarded as mere excursions (Schmid 1994).*

The tourism industry is gaining increasing attention as a major growth sector with great potential. As the largest contributor to the global economy, tourism employs 255 million people world-wide which accounts for 11% of the global workforce (WTTC 1996). It is estimated that tourism will sustain a 5% growth rate over the next decade resulting in an estimated 385 million jobs by the year 2006 (WTTC 1996). The tourism industry produces 10.7% of the world's gross national product (GNP) which is estimated to increase to 11.5% by the year 2006 (WTTC 1996). In 1995, international tourism receipts rose by 7.2% to almost US$372 billion (WTTC 1996). This increase follows a relatively constant trend established since 1985. Despite the recession in other industries the tourism industry is still growing at a faster rate than international economic growth (WTO 1993). Between 1970 and 1990 tourism grew by nearly 300% and it is expected to grow by half as much again by the end of the century. With such economic growth predicted, tourism is set to enjoy a further increase and the future long-term growth of the tourism industry, from a demand point of view, looks promising.

Despite the importance of the tourism industry to the global economy it is still very dispersed in its structure, comprising of many dispersed groups and services. In fact, tourism is probably the ultimately dispersed industry (Archdale 1993), tourism's relatively low degree of integration, further reinforcing the fragmented nature of the tourism industry (Go 1992). Furthermore, there are several changes in the tourism demand which are causing the tourism product to become increasingly difficult to control and manage. Firstly, potential customers are becoming more discerning when choosing their tourism destination. They require better value for money from tourism providers. Secondly, customers are travelling more frequently and to far greater distances throughout the world, largely due to the advances in airline and communication technologies. Thirdly, tourists are becoming more knowledgeable and adventurous in their choice of tourism destination (Collier 1989). As such, customers are becoming more global, which means they travel on a worldwide basis beyond normal travel patterns (Gamble 1991). And finally, customers are demanding a more individualised service and not those packaged by travel agents or tour operators (Collier 1989).

The development of new transport technologies has allowed us to travel further, faster. Each new technological breakthrough has allowed us to make another leap forward in our travel patterns (Hobson & Usyal 1993). However, it is already clear that the world's tourism infrastructure is being increasingly stretched and overloaded. The scale of growth in tourism will place great pressure on the facilities and amenities required by tourists (Cleverdon 1992). It is clear that future infrastructure projects will be developed at a slower rate than in the past. It will also not be at a pace that will keep up with the projected growth of tourism. Therefore, looking into the future, the growth of the tourism industry is likely to face more constraints from the supply side than the demand side (Williams & Hobson 1994).

4.3. Information Technology in the Tourism Industry.

Tourism is a very information intensive activity. In few other areas of activity are the generation, gathering, processing, application and communication of information as important for day-to-day operations as they are for the tourism industry (Buhalis 1994). Unlike durable and industrial goods, the intangible tourism product cannot be displayed or inspected at the point of sale before it is purchased. Furthermore, the tourism product is normally bought long before the time of use and away from the place of consumption. Information is the way the product is presented to the potential tourist. Consequently, tourists require a wide variety of specific information on the area, accessibility, facilities, attractions, and activities at the destinations as the provision of timely and accurate information relevant to the potential tourist's needs is often the key to successful satisfaction of the tourism demand (Williams 1993).

Therefore, largely due to the amount of information necessary the tourism industry has been significantly influenced by advances in Information Technology (IT) (Bennett & Radburn 1991). The original tourism applications were designed to make improvements in clerical and administrative efficiency, and as Gamble (1984) noted, 'the only difference being that computers work at electronic speeds'. Despite the technology playing an essentially passive role, operators gained substantial improvements in efficiency and in many cases this resulted in competitive advantage over other tourism providers. A new generation of systems has evolved, which are now playing a more active role in operations and management of the tourism industry. Broadly, current applications of computer technology in the tourism industry can be grouped into three main areas, operational, guest services and management information (Kasavana & Cahill 1987). IT is permeating all sectors of the tourism industry to varying degrees. The tourism product is becoming more dependent on IT, and it is predicted to be more so in the future. More cross-sector alliances, direct links between the provider and the tourist and the drive towards efficiency will require access to the processing and translation of large amounts of information. This information will need to be processed faster, more reliably and in more understandable format to enable prompt decision-making.

The tourism product has many interrelated but separate components: market demand; travel to and from the destination; destination services, attractions and facilities; and the marketing of the product (Mill & Morrison 1985). To these components a fifth could be added, the management of the tourism product, since the interrelationship between the four traditional components are becoming increasingly complex and IT is playing a more executive role in their management.

Tourism Market Demand: The key to successfully satisfying tourism demand is the provision of information to the customer. This information must be timely, accurate and relevant to the customer's needs. In tourism, the product is largely intangible, perishable, heterogeneous and volatile, and as such, it is the information provided to the potential tourist that is recognised as being the product (Bennett & Radburn 1991). Potential tourist, therefore, rely on a wealth of information before making a decision.

In providing this information, there is a series of interdependencies evident. The travel agent is often the key link between the tourism provider and the potential tourist. The potential tourist relies on the travel agent to provide the broad range of information at the point of sale, and the tourism provider relies on the travel agent to provide the desired information to the tourist. Travel agents must, therefore, have instant access to the information required in order to satisfy tourist demand. To aid them in this task travel agents are installing central reservation system (CRS) terminals of the airlines who have realised the potential of the powerful travel agency network for distributing its product through CRS. Although originally developed to control reservations inventory and efficiently manage airline scheduling, the CRS industry is now much more complex with links to other sectors of the tourism industry, such as hotels, car hire companies and tourist organisations.

Traveling To and From a Destination: Travelling to and from the destination is the second element of the tourism product and is often the largest proportion of the expenditure in the purchase of the whole tourism product (Hitchins 1991). IT investment has been significant in terms of organisation and scheduling of the trip, but minimal in terms of improving the travel experience in its own right.

The Tourism Destination: The third component in the tourism product is the tourism destination. The tourism destination uses IT to plan, manage and service its tourism product at the destination. The future success of the destination is dependent on the utilisation of IT. The main focus of IT in this area has been in the improvement of the tangible facilities offered.

Marketing of the Tourism Product: Marketing of the tourism product has received the most attention in terms of investment and research. In the tourism industry, potential tourists are unable to sample the product prior to purchase and must, therefore, learn about the product through the information provided.

The better the quality of this information the more likely the potential tourist will be to formulate a realistic impression. In an increasingly global market, both tourism providers and potential tourists are dependent on access to good information. Inevitably, the effective marketing of tourism is becoming increasingly dependent on IT. Marketing distribution channels which exist now, and in the future, can not survive without investment in IT.

Management of the Tourism Product: With many different sectors and cross-sector alliances forming, a very tight control of the tourism product is required. Computers and other communication technologies are assisting management to coordinate and control the many inter-dependent sectors of the tourism product. Management and coordination of the tourism product through increased investment in IT is improving the efficiency in many sectors of the tourism industry.

4.4. Applications of Virtual Reality in the Tourism Industry.

The tourism industry has already benefited significantly from the implementation of IT. IT systems govern many aspects of the tourism product from airline and hotel reservations to multimedia packages designed to market tourism destinations. VR adds another dimension to the marketing of the tourism product. Features of the tourism industry lend themselves readily to VR (Cheong 1995) and the technology creates a vastly different approach to the representation of information. With VR, participants find themselves in the same dimension as, and are immersed within, the information. The Virtual Environment (VE) may be augmented by sensory feedback of sight, sound, and even touch. This provides an excellent mechanism to access, conceptualise, and manipulate tourism information.

The relationship between tourism and technology has been rapidly evolving. The use of computers within the tourism industry has evolved, in a very short period of time, from a simple back-office system to an indispensable management, marketing and financial tool. With the introduction of VR this relationship has moved one step further to being a tool capable of helping both management and customers to make well-informed decisions. In the field of tourism, images have been used to promote destinations for a considerable time, either in the form of video films (Hanefors & Larson 1993) or more recently electronic brochures (Archdale 1993). Coates (1992) imagines travel agents taking their customers on VR tours to any part of the world, exploring in depth the experience before the real trip and reviewing it afterwards. VR has been identified as one of the most exciting technological developments of the future, where different experiences can be created through the use of computer-mediated imaging and graphical environments. Potentially, there is an infinite number of tourism experiences which VR can simulate. The benefits engendered by VR applications at present manifest themselves in two distinct levels:

1. Macro Level - concerning tourism policy and planning.
2. Micro Level - concerning the provision of information and the marketing of the tourism product.

4.4.1. Tourism Policy and Planning.

The process of developing a tourist destination is a long and tedious one. Tourism master plans encompass numerous detailed physical layouts of the destination together with the proposed tourism infrastructure such as roads, hotels, restaurants, and other service facilities. These amenities are often depicted on paper as 2D representations that do not facilitate the tourism planners' conceptualisation of the destination (Cheong 1995). Although digitised representations of the proposed site can be subjected to numerous operations, such as image rotation and enlargement, by CAD programs to achieve conceptual versatility, they are still perceived as 2D (Cheong 1995).

When planning a proposed tourism destination, a 3D representation is necessary in order to obtain a correct perspective of the site. Undesirable elements of the destination's design may not be obvious from 2D representations and, therefore, may be overlooked in the design of such a location. These limitations and drawbacks can be alleviated with the use of VR. All relevant information with regards to the proposed tourism destination can be incorporated into the computer system. Utilising information such as the geographical layout of the site, the infrastructure and tourist activities, a designer can generate a Virtual tourist destination which tourism planners can enter and, more importantly, interact with in real-time. With these tools planners can analyse the layout of the proposed infrastructure. Each destination can, therefore, be changed or adapted in order to determine the most appropriate layout for the proposed destination. With this approach to tourism destination development previously unforeseen discrepancies can now be discovered prior to any real expenditure. Therefore, VR has the potential to serve as an invaluable tool in the formulation of tourism policy and in the tourism planning process.

4.4.2. Information Provision and Marketing of the Tourism Product.

Traditional methods of tourist information could be enhanced by VR. VR can serve as an invaluable method of providing necessary information to potential tourists and, therefore, act as a marketing tool for the tourism industry. VR systems have the ability to provide potential tourists with the opportunity to experience destinations and their respective attractions and facilities. Traditional sources of tourism information only provide potential tourists with short and often rather limited glimpses of tourism destinations that may be inadequate to enable them to make informed decisions. The underlying significance of VR is its ability to 'bring the experience to the customer' and this virtual experience of the destination will further increase the customer's desire to actually visit the location (Coates 1995). The traditional methods of tourism information possess no involvement on the part of the potential tourist. This, to an extent, limits their effectiveness as a means of encouraging a potential tourist to a particular destination.

Most of the exciting developments in destination and product presentation, such as multimedia and VR, are particularly relevant at the information stage, within the search process for a suitable holiday destination as described in Goodall (1990) as 'pull' factors (Riley & Van Doren 1992). It is in the area of the provision of information and the marketing of the tourism destination that VR is most useful at present. The application of VR described in this thesis is in the provision of a Tourism Information

System (TIS). Such a system includes the traditional types of tourism information, images, text, sound, animation and video, but also incorporates a VE which allows the potential tourist the opportunity to walkthrough the tourism destination at their own leisure as well as interact with their surroundings. This VE is created to help tourists to generate realistic impressions and expectations of what to expect at such a destination and, thus, provide them with extra information to make an informed decision on their tourism destination.

4.5. Preliminary Design Considerations.

This section addresses issues associated with the design of a VR application to the tourism industry. The application that has been developed is called Desktop Virtual Reality Version 3 (D-VR3). D-VR3 is a Desktop VR application which is designed to be a TIS which provides the user with all the traditional types of tourist information along with allowing the user to participate in an interactive Virtual Walkthrough of tourism locations as they appeared in the past, as they are at present, or as they may be in the future. These Walkthroughs are not intended to replace the experience of visiting the original location but instead are merely used to supplement them for information purposes. Text, oral, pictorial and video presentations, the traditional types of tourism information, address subsets of human capacity. In contrast, Virtual Environments (VEs) provide a context that includes the multiple nature of human intelligence: verbal/linguistic, logical/mathematical, auditory, spatial, inter-personal and intra-personal. Adding VEs to the traditional types of tourist information will increase the user's knowledge base and, therefore, aid them in making a conscious, well informed tourism decision.

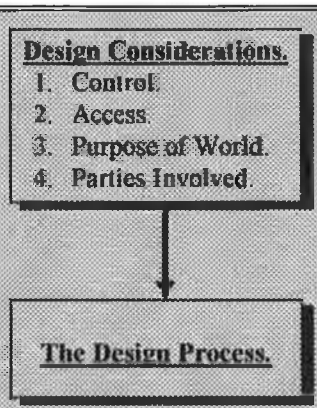

Figure 4.1. Preliminary Design Considerations.

The design process begins with defining the parameters of a VE. This is in essence, the workspace that a VE designer has to deal with - like a painter's blank canvas. Before dealing with the intricacies of Virtual World design it must be understood that certain characteristics impose design considerations on a VR system. The design considerations, illustrated in Figure 4.1., are an extremely important part of the design process as they are the factors which quite often dictate the type of VE created, the software used to create that environment and the interface used to interact with that environment. The accessibility of affordable systems, the ease with which people can use them and the degree of control that individuals have over their participation in VR will all be influenced by designers and shall now be dealt with in some detail.

4.5.1. Power and Control.

There is little doubt that VR is an empowering technology; but exactly who is being empowered? This is, in part, a design consideration. The current development of relatively inexpensive systems along with high-end models indicates that the technology will soon be widely available. But once a participant enters a Virtual World who is in control? This is a very important matter that should be dealt with in great detail before building the VE. There are three types of control:

Application Control: The design of VEs can induce passivity or facilitate creativity. An application that restricts the interactivity of the medium usurps the participant's power and, therefore, the potential of the application. It is the aim of D-VR3 to provide an interactive experience and to allow the participant as much control as is possible over his activities while in the VE.

Social Control: When VR becomes commonplace, how will the power of the technology be distributed? The applications of VR to the tourism industry are stunning in concept, but tourists may respond to this potential in two ways. One response focuses on the new experiences it will afford tourists, the other centres on the additional control that may be held over tourist's behaviour. It is the aim of D-VR3 to provide the tourist with as much social control as possible in order to allow them to gather the necessary information to make a conscious well informed tourism decision.

Personal Control: Once a person enters a VE he must have control over his own activities, in other words, his "Virtual Body" must be controlled by him alone. In many VR systems there is a mechanism to allow a designer to take control of the participants movements and to lead them around a world. This mechanism leads to confusion over who exactly is in control and disorientation about where they are within the environment. For these reasons such a mechanism has not been included into the Virtual Worlds in D-VR3.

4.5.2. Access.

A decision about a participant's access within a VE is a very important issue which needs to be addressed prior to constructing VEs as it not only dictates who is allowed to enter the environment but also which areas they are allowed to access once they have entered the environment. In relation to D-VR3, there are four issues with regard to access:

Who Should be Allowed Access? D-VR3 is intended to be an open VE. Therefore, anyone with the correct hardware and software should be granted access to the VE.

Which Areas Should a Participant be Allowed to Access? There are two types of access system currently available to designers of VEs and they are:

1. Branching Systems - A branching system constrains a user to predefined pathways like paths, roads or sidewalks within the Virtual World.

2. Free-Roaming Systems - A free-roaming system permits the user to move anywhere within the Virtual World.

D-VR3 is designed to be a free-roaming system as this type of system allows the participant maximum versatility within the VE and, thus, increases the potential of the application.

The adoption of a free-roaming system approach requires flexible mechanisms through which a participant in one world may access another world. A common solution to this problem is through the use of a portal approach. A portal approach attempts to simulate fluid movement between worlds. This is achieved by placing a portal, an entry or exit point, in each of the worlds, Figure 4.2. Once a participant moves into the portal in the first world he leaves that world and is immediately transported to the location of the second portal in the second world. In order to avoid this unnatural form of movement, a designer should endeavour to create structures that make travel between the worlds appear more realistic. In the case of D-VR3 all portals are placed in doorways, thus creating the illusion of entering or exiting an environment. The use of a portal approach in D-VR3 becomes an essential part of the Virtual Experience as it allows large worlds to be divided up into smaller worlds, thus, allowing real-time interaction since in some VR engines, the world is redrawn very slowly if the world is large.

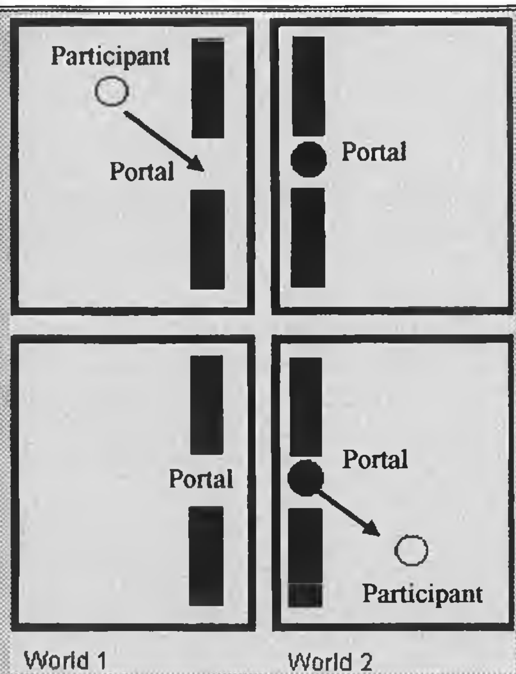

Figure 4.2. The Portal Approach - Movement between Virtual Worlds.

How Many People Should be Allowed Access to the World at One Particular Time? It is clearly desirable that a VE should support multi-user operation. Such a facility requires the use of multi-user domains and, while incorporated within D-VR3, performance is host system dependent. Further consideration of this issue is provided in Chapter 5.

How Should Access to a Virtual Environment be Provided Using D-VR3? To facilitate the participant's initial introduction into the Virtual Worlds, a multimedia-based environment is adopted. This multimedia-based environment, described in more depth in Chapter 5, is used to provide the tourist with access to each environment in an easy to use informative way. In this multimedia environment the traditional types of tourist information are provided including text, pictures, sound and video. All of these traditional types of tourist information are passive media and the tourist is led on a tour of the information. D-VR3 goes one step further and includes Virtual Worlds that a tourist can actively explore and, therefore, actually participate in the tourist information process.

4.5.3. The Purpose of the Virtual Environment.

In the case of D-VR3 the aim is to provide a potential tourist with the opportunity to experience a multi-participant, realistic, interactive and real-time walkthrough of real-life tourist destinations on commercially available PCs in order to provide the tourist with enough information to make a conscious, well-informed tourism decision.

4.5.4. The Parties Involved in Virtual Environment Creation.

The creation of VEs can be a very intricate procedure. Primarily the procedure is based on an interaction between the three parties involved - the engineer, the designer and the participant. Hardware and software engineers create the interactive devices and tools; they focus on how systems work. Designers work with engineers to ergonomically refine prototypes and with participants to tailor systems to individual purpose; they focus on the people using the technology. Participants are concerned with costs and benefits of the system, and focus on their particular VE. Because engineers are primarily concerned with implementation, designers with evaluation, and participants with function, they have different conceptual models of the system. To create effective VEs requires good communication between these parties and it is essential that each party must have an intimate knowledge of the other party's role in VE creation, Figure 4.3.

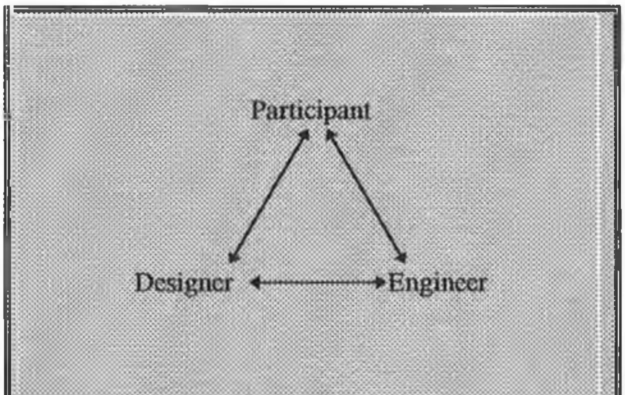

Figure 4.3. Design Relationship

The engineer's model of the system involves the functional intricacies of hardware and software, and therefore, the engineer's priorities are structured around what is technically possible, or at least feasible. In the construction of the Virtual Worlds in D-VR3 the role of the engineer is not as significant as that of the participant and the designer. This is because it is not the aim of the project to design tools for the creation of VEs but instead it is the aim to adapt and apply existing VR tools in order to create worlds suitable for the tourism industry.

Most people are more interested in what they do in a VE than how the system works. The participant's model of the system is based on personal experience with it. From the participant's perspective within the VE there are several considerations that need to be addressed:

Virtual Position: A participant's primary understanding of the Virtual World, and his position within it, is derived from what is seen and heard when entering the world. It is essential that the correct information regarding the participant's position within the world is presented to the participant.

Virtual Body: In a VE a participant does not necessarily need a body; he can exist as a floating point within the world. The "Virtual Body" can take up any guise. In the case of Virtual Worlds in D-VR3, the appearance of the virtual self depends entirely on the engine that one is using as some of the packages use floating points while others use representations of a human body.

Virtual Actions: There are many different operations that a participant may undertake while present within a VE. These operations may be classified within the following four categories of behaviour:

Relocation: A participant may move around a VE in a manner similar to the way he does in the physical world, by walking and turning, bending and reaching. A participant may also fly smoothly with variable speed in any direction using a hand gesture, joystick or trackball. In VR, the concept of "distance" is optional and relocation is independent of time and space.

Manipulation: A participant may move virtual objects in a variety of different ways and with a variety of different devices. A participant may cause an object to move, or patterns of movement for an object can be programmed.

Construction: Presently, Virtual Worlds are built from the outside, and are then entered. The ability to interactively create and alter these environments is being developed in the form of software toolkits. The ability to shape and visualise information depends on the tools that are used.

Navigation: Navigating around a small Virtual World is not difficult. However, where a world is complex or extensive, or when connectivity between a number of worlds is required, difficulty significantly increases and new methods of locating objects, and navigating between domains may be needed.

Virtual Companions: The feasibility of sharing virtual spaces has been indicated. Yet some questions remain unanswered; how will other people be represented in a Virtual World; does each participant control his appearance or does he convey a fixed form? The usual concept of "who" people are implies some consistency, but in VR people can adopt an ambiguous form or multiple personae. The meaning of personal appearance changes; in the physical world certain behaviours may be predicted from the way a person looks.

Designers focus on the way people access and interact with VEs. The designer's objective is the creation of comfortable, functional worlds that satisfy the needs and intentions of the participant. There are four aspects of this design task:

1. Designers work with engineers to tailor technology to people's physical and psychological needs by testing the usability of systems and suggesting refinements. In the design process an overlap occurs between the designer and the engineer when choosing the VR system and software and hardware. In this application there is no interaction between the engineer and designer but in order to combat this the different types of VR systems and VR hardware and software have been analysed in order to select the most appropriate for the tourism industry and for the specific application.

2. Designers work with participants to customise Virtual Worlds. This requires a dual awareness of the individual needs and preferences of the participant and of the capabilities of the technology. This is an integral part of any design process and during the design of D-VR3 liaison occurred between the designer and the participant at every given opportunity. The application has been presented to staff from Siemens Nixdorf, the project's Industrial sponsors, staff from the interpretative centres in Glendalough and Newgrange, architects from the Office of Public Works (OPW), tourists visiting the Monastic Enclosure at Glendalough and colleagues at the Dublin Institute of Technology (DIT), Cathal Brugha St. This exercise has been invaluable in providing information on the appearance and behaviour of objects within a VE and on different interface devices necessary to interact with that environment.

3. Designers compose prototype worlds that contain graphical contexts and interactive possibilities appropriate to particular applications. A prototype world is composed of the Virtual World's appearance and behaviour, appropriate libraries of objects and sounds, and the vocabulary of interactions available in the virtual world.

In the case of D-VR3, the Computer Section at D.I.T Cathal Brugha St. has been chosen as the prototype world. The reasons for these choices are as follows. Firstly, it is a familiar structure and, thus, little time would be spent on familiarisation with the intricacies of its design. Secondly, the scenario of multiple participants could be easily examined via the college's computer network. And finally, because D-VR3 could be experienced in the Computer Section it would be relatively easy to compare the prototype model to the actual building.

4. Designers evaluate effective worlds by observing the learning and performance of participants. The ideal Virtual World is fully explorable, thus allowing the participant to build a specific cognitive model through experience. A Virtual World is inherently information-rich, and if feedback is explicit, the discovery process appears to be a relatively efficient way to learn. This is dealt with in more detail in Chapter 6.

4.6. Summary.

In this chapter the importance of tourism to the global economy, its growth rate, and its dependence on information have been discussed. The tourism industry has many interrelated but separate components and these components which becoming increasingly complex have to be managed effectively (Williams 1993). IT is playing an ever-increasing role in this management.

Tourists are dependent on accurate, relevant and timely information in order to aid them in their travel decisions. In addition, the tourism product is largely intangible, perishable, heterogeneous and volatile, and it is, therefore, the information provided to the potential tourist that is recognised as being the product. The provision of some tangible elements, such as video clips, animation and virtual walkthroughs of tourism destinations, hotels attractions and local environment can reduce some of the intangibility of the tourism product, especially for destination-naïve tourists. For this reason, VR technology is likely to have a major impact on the future of the tourism industry.

The uses of VR in the tourism industry were also discussed and the design considerations for a specific application to the industry, D-VR3, have been established. D-VR3 is a PC-based application that provides tourists with all the traditional types of tourist information in addition to highly vivid VEs, with which they can interact in real-time and experience the amenities that the destination has to offer. Using such a tool, potential tourists may be provided with sufficient information to enable them to formulate realistic expectations of intended destinations. The next chapter describes the design process used in the construction of VEs present in D-VR3.

Chapter Five

The Design Process - The Making of D-VR3

5.1. Introduction.

A Virtual Environment (VE) is a world of pure information that one can see, hear, and touch. The world is carefully adapted to human activity so that one can behave naturally in it. Any imaginable environment and experience may be created using VR. A VE may be informative, useful, and fun; it can also be boring and uncomfortable. The difference is in the design.

This chapter addresses issues associated with the design of D-VR3, an application of VR to the tourism industry. The aim of this application is to provide a potential tourist with the opportunity to experience a multi-participant, realistic, interactive and real-time walkthrough of real-life tourist destinations on commercially available PCs in order to provide the tourist with enough information to make a conscious, well-informed tourism decision.

D-VR3 is a real-time graphics development environment for building VR applications. Although the main focus is on the Windows 95 version of D-VR3, it must be understood that D-VR3 is also available on PCs running Windows 3.1 or DOS to a limited degree. D-VR3 contains routines to control, interact, view, and change objects in a VE. The user controls any of a number of input devices allowing movement and manipulation of objects in the environment. D-VR3 automatically displays a view into the environment and controls all aspects of the display, including the view perspective, shading, texture mapping, display updating, and querying input devices. For these reasons D-VR3 is a flexible and useful application.

As a research vehicle, D-VR3's primary aim emphasises functionality at the expense of performance. Premature optimisation is the most common source of difficulty in software research (Bricken 1991). An application is best improved after it is functioning properly. The main aim in D-VR3 is to demonstrate that a function can be carried out at all rather than demonstrating how well it can be conducted. D-VR3 is intended to show the way forward, and is, therefore, constantly undergoing revision and iterative refinement.

In summary, D-VR3 is a significant effort to provide Desktop VR, process, and communications management for arbitrary sensor suites, software resources, and Virtual World designs. As such, D-VR3 provides a strong integration environment for any team wishing to construct, experiment with, and extend VR systems.

5.2. The Design Process - Virtual World Creation.

At this juncture it is necessary to make two points. Firstly, Virtual World creation of a tourism location at present should be undertaken using Desktop systems. It is simply not practical to use immersion systems, because the headset cannot be endured for long enough periods for the application to be useful. Secondly, Virtual World creation should not be restricted to those with computer programming expertise. Designers, graphic artists, and indeed anyone involved in the preparation of presentation media should be able to create Virtual Worlds. VR systems, like D-VR3, need a "front end", an interactive utility that allows the simple creation of Virtual Worlds, without the need for programming skills. Even programmers prefer using an interactive utility because world creation becomes a considerably faster process. This aspect is key to popularising the use of VR systems.

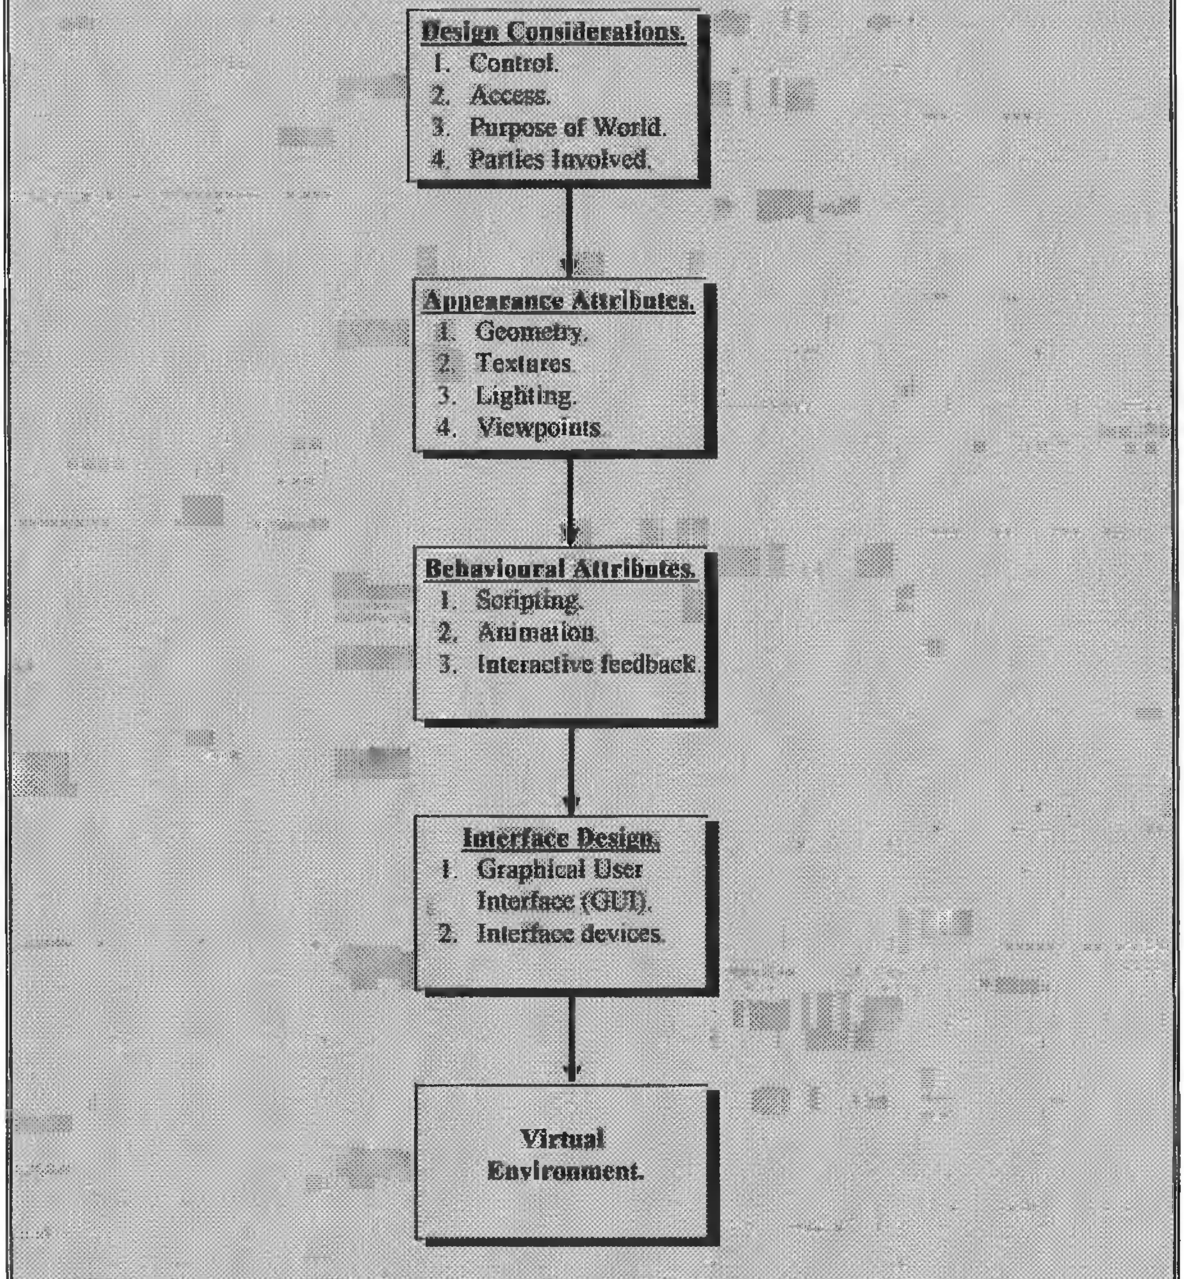

Figure 5.1. The Stages Involved in the Design of a Virtual Environment.

Independent of the tools used to create a VE, the processes that a designer must undertake when creating such an environment are always the same. There are many stages involved in the creation of a VE, shown in Figure 5.1., but the creation of appearance and behaviour are arguably the two most important. The design process begins with defining the parameters of a VE. The next step is to add contents to an environment. The basic graphical entity in a environment is an "object." Objects can be constructed using a variety of modelling techniques. These techniques result in a mathematical description of an object that has a surface composed of polygons, simple geometric shapes, that can describe any surface. Once the objects are defined attributes such as shading, light and textures are applied to them. In the final stage of Virtual World creation behavioural attributes can then be applied to certain objects within the world. Once the VE is complete an interface device has then to be decided upon.

Before dealing with the intricacies of Virtual World design it must be understood that certain characteristics impose design considerations on a VR system. Since these considerations have already been addressed in chapter 4, and the parameters set, it is appropriate to progress to the next stage of the design process, the creation of appearance.

5.3. The Creation of Appearance.

The creation of appearance begins with 2D shapes such as polygons, circles, curves, and typographic elements. These 2D shapes are then extended into a third dimension through a variety of techniques such as extrusion, surfaces of revolution, and lofting. Once the coordinate geometry for a VE has been designed, a rigid skeleton has been established. If one was to view the environment at this stage if would look like a line drawing or, in computer graphic parlance, a wireframe. The next step in the creation of appearance is describing the colour, shading, texture, lighting and shadows which are applied to the geometric skeleton. Appearance may be described in many ways. Often computer graphics designers create surface appearance by "painting" a 2D image and then wrapping it around a 3D object. This process is known as texture mapping. Alternatively, designers may use mathematical models that generate desired patterns such as wood grain or marble. Different kinds of material such as metal and glass are simulated using other mathematical models.

Surfaces are unfinished unless they are properly illuminated. The most exquisite museum exhibit is not shown to best advantage with the drapes drawn and the light off. Lighting has the capability to shown a VE to best advantage. Lighting parameters include light colour, intensity, placement and shadow. A designer can specify ambient light and directional light sources, which have direction and brightness. In all of the Virtual Worlds in D-VR3 texture mapping, lighting and shading of each polygon is automatically handled, and the participant is always presented with a realistically shaded view.

Typically, the starting point in the creation of appearance for any VE is the description of the basic geometry of the world. Therefore, it is fitting that the first topic to be discussed is the description of the world's geometry.

5.3.1. Geometry.

The basic graphical entity in a Virtual World is an "object". An object is described as a collection of polygons each of which is described by a collection of vertices, or points. For example, the simplest polygon is a triangle that consists of three vertices. An object is also described by its appearance: its given attributes such as colour, texture, lighting and size. In D-VR3, all objects are organisationally identical. Only their structure, or internal detail, differs. This means that a designer needs only one metaphor, the object, for developing all aspects of the world. The way that the object is viewed and heard can be ignored for the purposes of the current discussion; one must deal with the appearance of the world and not worry too much about its behaviour. Suffice to say that each object interacts with another object in some way. This is simulated by each object opening a link with another object. This link is a communications channel, through which information may be passed between objects. This link is actually a logical extension of a certain component of an object, discussed later in section 5.4.

Objects in the Virtual World can have geometry, hierarchy, and other attributes. These attributes can be added to the system without requiring changes to the object data structures. The capabilities of objects have a tremendous impact on the structure and design of a VE. These attributes are now discussed.

Position/Orientation: Like a Virtual Body, all objects are positionable and orientable. That is, they have a location and orientation in space. Most objects can have these attributes modified by applying translation and rotation operations. These operations are often implemented using methods from vector and matrix algebra.

Hierarchy: An object may be part of an object hierarchy with a parent, sibling, and child objects. Therefore, an object would inherit the transformations applied to its parent object and pass these on to its siblings and children. Hierarchies are used to create jointed figures such as robots and animals. They can also be used to model other things like the sun, planets and moons in a solar system.

Object Geometry: The modelling of object shape and geometry is a large and diverse field. Some approaches seek to very carefully model the exact geometry of real world objects. Other methods seek to create simplified representations. Most VR systems sacrifice detail and exactness for simplicity for the sake of rendering speed. This, of course depends on the purpose that the world is designed for. The purpose of the world designed in D-VR3 is to provide the participant with a perception of the tourist site that is modelled. This is achieved by presenting him with a photorealistic world that he can walk around and explore at his own leisure.

The simplest objects are single dimensional points. Next come the 2D vectors. Many CAD systems create and exchange data as 2D views. This information is not very useful for VR systems, except for display on a 2D surface within the Virtual World. There are some programs that can reconstruct a 3D model of an object, given a number of 2D views. The different types of geometric modelling methods are discussed in more detail in Appendix B. The choice of method used is closely tied to the rendering process used. Some renderers can handle multiple types of model, but most use only one, especially for VR use. The modelling complexity is generally inversely proportional to the rendering speed. As the model gets more complex and detailed, the frame rate drops.

In Appendix B the different type of techniques available for object construction are addressed. The actual techniques that were used for object modelling in the Virtual Worlds in D-VR3 depend on the actual software packages used and are discussed in sections 5.11.1., 5.12.1., and 5.13.1. The next stage that is addressed in the creation of appearance is the inclusion of shading and lighting.

5.3.2. Shading and Lighting.

The description of the Virtual World thus far has included objects represented by wireframe drawings. Improving the appearance of the Virtual World includes adding certain attributes to these wireframe drawings in order to convert them to more solid objects. These attributes include adding shading and light to these objects.

Creating a shaded object involves describing the object a little differently. Instead of describing the object in terms of points and lines, in a shaded situation the object is described in terms of flat faces. A face is simply a flat part of the surface of an object, and it can be defined by listing the vertices around its perimeter, Figure 5.2. In the figure below there is one face that is shaded and it can be defined by listing the vertices in an anti-clockwise fashion. In the figure below the shaded face can be described as 1, 2, 6 and 5.

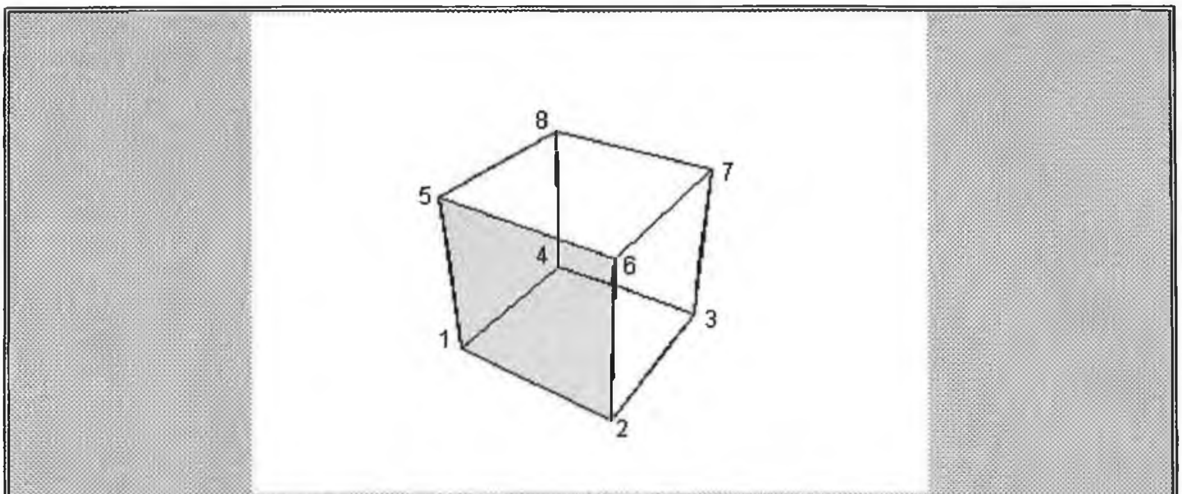

Figure 5.2. A Wireframe Diagram of a Cube with one Face Shaded.

In many graphics tools faces can only be viewed from one side. This is mainly due to do with saving drawing time by not drawing the inside of the object. The order in which the vertices are listed indicates the side of the face that is visible. A face does not necessarily have to be filled with a single colour. By filling each face a pixel at a time a number of shading methods are possible. The simplest and fastest shading method is called flat shading, in which the computer assigns a single colour to each polygon without textures. Gouraud shading, the next step up, is slower but can smooth the colouring of the surface across the polygons and apply texturing information, resulting in a more realistic look. This is the method typically used in interactive renderers. Phong shading is slower than Gouraud shading because it represents the lighting in a scene more accurately, producing effects such as the bright highlighting, but not reflection, found on plastic or metallic surfaces. Finally, ray tracing is an extremely slow shading method that produces photograph-quality results. This method can simulate reflective (chrome), refractive (glass) and transparent (coloured glass) surface qualities that contribute to the photographic illusion. Most artists use Phong shading or ray tracing to produce cinematic animation frames.

So far in this section the topic of filling in the faces of an object has been addressed. If a designer wishes to illuminate objects in a Virtual World, then the world must have the ability to fill faces with different shades of the same colour. One possible combination is a 256 colour display with 16 colours and 16 different levels of brightness each. Any display with less than 256 colours will not provide enough colours for the different levels of brightness.

In a Virtual World, lights can be divided into two categories: ambient or directional. The easiest light source to include in a VE is an ambient light as it provides an equal amount of light to all parts of the world. Directional lights, on the other hand, have position and may have orientation, colour, intensity and a cone of illumination. Directional light only illuminates an object from one side. Each object face reflects a certain amount of directional light. This will depend on the object's angle to the light: the more it faces towards the light the more light it reflects. One type of directional light source is a spot-light, which has a direction, but no location, similar to the light cast from the sun. A second type of directional light would be a point-light source. The more complex the light source, the more computation is required to simulate its effect on objects.

5.3.3. Textures.

Even with the most advanced lighting and shading model, a Virtual World would still look relatively bland with the surface of objects consisting of only a few colours. Texture-mapping, the application of photographic and synthesised images onto the surface of a simulated object, is the key to enhancing the realism of any VR application. A photographic image can be imported and mapped onto a wall, for instance, and one instantly has an environment that the brain is more than willing to believe in, even without full "photorealism" effects (Stoppi 1992). Texture mapping requires that the source 2D textures be 'wrapped' onto the multiple surfaces of the objects in the world.

Many studies have shown the perceptual benefits of using textures compared to just using shading techniques in simulations (Coull 1991). Textures make objects appear more life-like, and less cartoon-like. Thus, they allow VR systems to present an image that is closer to that of the real world.

Textures also provide addition depth cues for object position. In general, it is easier for a participant to determine an object's distance in a textured 3D world than in a non-textured world. With textures, depth cues are derived from an object's edges and surfaces - in the case of a flat shaded object, they are only derived from the object's edges. The additional depth cues enhance the viewers sense of 3D perspective and allow easier interaction with objects.

5.3.4. Other Considerations.

Many other considerations exist when creating the appearance of a Virtual World. The most important of these other considerations is a participant's starting position and orientation. This is the participant's position and the way that they are facing once they enter the Virtual World. This is important. A designer must place a participant's starting position in close proximity to the main objects in the world and they must be facing an appropriate direction in such a way as to gain a good impression of the world as soon as he enters it.

Another consideration is where to place a number of multi-user starting positions. The same rules apply to multi-users' start positions and orientations as apply to a single participant's start position and orientation. But other considerations may also apply. For each user that the world is programmed to cater for one user position has to be placed in the world. Therefore, if the system is programmed to cater for 20 users there must be 20 user start positions. The world can also be created in such a way that the users do not all start from the same location in the Virtual World. This may be conducted to reduce the confusion if a large number of users enter the world at the same time. Another consideration is the user's representation in the multi-user world. This concept is known as user embodiment. In some VR applications user embodiment is extremely important while in others it is not so important. A user may be represented by anything from a floating point to a dynamic computer generated representation of themselves.

As mentioned earlier, the two most important factors that must be considered when creating a VE are the creation of appearance and the creation of behaviour. Therefore, the next topic to be discussed is the creation of behaviour in the VE.

5.4. The Creation of Behaviour.

VR is not only about realistic sensory substitution; it is also about realistic behaviour. Behaviour may be described as an object's reaction or reactions to events. Trying to model complex interactions between all the objects in a Virtual World may be impossible if a designer tries to tackle all possible interactions at once. For this reason, most worlds allow a form of object-oriented programming, where each object in the world is provided with its very own program that will define its reactions to certain interactions. By breaking the interactions down into simple rules for each object some very complex interactions are possible. The rules that are associated with an object define its personality. These rules include gravity and weight, elasticity, command structures, collision detection for manipulating objects, movement constraints, animations and auditory feedback. Portals also add to the interactivity quotient. The more rules that are applied to each object the more complex the resulting behaviour. The main aim of adding behavioural attributes to an object is to make the Virtual World appear more realistic and, therefore, more engaging. By choosing the behavioural attributes carefully the Virtual World can become more natural and more useful.

A Virtual World consisting purely of static objects is only of mild interest. This, of course, depends on the purpose for which the world was created. However, interaction is a key component of any Virtual World. This requires some means of defining the actions that objects take on their own and when the user, or other objects, interact with them. This is referred to generically as World Scripting. Scripting is divided into six basic types and each type is outlined.

Motion Scripts: Motion scripts modify the position, orientation or other attributes of an object, light or camera based on the current system tick. A tick is one advancement of the simulation clock. Generally, this is equivalent to a single frame of visual animation. For simplicity and speed, only one motion script should be active for an object at any one instant. Motion scripting is a potentially powerful feature, depending on how complex one allows these scripts to become. Care must be exercised since the interpretation of these scripts will require time, which will cause delays in frame rates.

Additionally, a script might be used to attach or detach an object from a hierarchy. For example, a script might attach the user to a car object when he/she wishes to drive around the Virtual World. Alternatively, the user might 'pick up' or attach an object to himself/herself.

Trigger Scripts: Trigger Scripts are invoked when some trigger event occurs, such as collision, proximity or selection. The VR system needs to evaluate the trigger parameters at each tick. For proximity detectors, this may be a simple distance check from the object to the participant. Collision detection is a more involved process. It is desirable but may not be achievable with certain VR tools.

Connection Scripts: Connection scripts control the connection of input and output devices to various objects. For example a connection script may be used to connect a mouse movement to a virtual hand object. The mouse movements and position information are used to control the position and actions of the hand object in the Virtual World. Some systems build this function directly into the program. Other systems are designed such that the VR program is almost entirely a connection script.

Procedural Modelling: A complex simulation could be used that models the interactions of the real world. This is referred to as procedural modelling. It can be a very complex and time-consuming application. The mathematics required to solve the physical interaction equations can be very complex. However, this method can provide a very realistic interaction mechanism.

Simple Animation: A simpler method of animation is to use basic formulas for the motion of objects. A very simple example would be "Rotate about Z axis once every 4 seconds". This might also be represented as "Rotate about Z 10 radians each frame".

A slightly more advanced method of animation is to provide a 'path' for the object with controls on its speed at various points. These controls are sometimes referred to as "slow in-out". They provide a much more realistic motion than simple linear motion. If the motion is fixed, some systems can pre-compute the motion and provide a 'channel' of data that is evaluated at each time instance. This may be a simple lookup table with exact values for each frame, or it may require some sort of simple interpolation.

Interaction Feedback: The participant must be given some indication of interaction feedback when the virtual cursor selects or touches an object. Crude systems have only the visual feedback of seeing the cursor (virtual hand) penetrate an object. The user can then grasp or otherwise select the object. The selected object is then highlighted in some manner. Alternatively, an audio signal could be generated to indicate a collision. Some systems use simple touch feedback, such as a vibration in the joystick, to indicate collision. Interaction feedback will depend on the VR system and the manipulation devices chosen.

So far in this chapter the processes to create a realistic interactive and real-time VE have been described. Then next step is to describe the actual steps taken to develop the Virtual Environments (VEs) used in D-VR3. In an attempt to describe the processes undertaken the process is divided it into two sections. The first section describes the steps taken to produce the prototype model of the Computer Section in Cathal Brugha St. The second section describes the production of the Actual Model of the Monastic Enclosure at Glendalough.

5.5. The Prototype - The Computer Section, D.I.T. Cathal Brugha St.

One important method for defining both the participant's and the designer's needs is the creation of a prototype. A prototype is a working model - a shell or skeleton of the final application. Prototypes contain outlines of the screens, menus and models that the application will use. They provide a mock-up of the visual elements and explain the processes required. Prototyping not only helps the participant visualise the final product, but they also help the designer to put together the component parts of the final application.

The actual techniques used to create Virtual Worlds in D-VR3 will be described with reference to the specific software packages used. Choosing the appropriate tools is an extremely important part of any design process and this is the main reason why so much emphasis has been placed on the prototype model in D-VR3. The first step in the choice of appropriate software entails a comprehensive analysis of the desktop VR packages available. In an attempt to choose the correct software for D-VR3 three different categories of software have been examined. These categories include:

1. Three-Dimensional Modelling and Visualisation Programs;
2. Games Engines;
3. Three-Dimensional Graphics and Animation Packages.

Therefore, in this section the techniques used in the creation of the prototype model will be described with special reference to these three types of software packages.

5.6. Three-Dimensional Modelling and Visualisation Packages.

The first category of package analysed was the commercially available three-dimensional modelling and visualisation programs. These packages were chosen because of their ability to design highly accurate Virtual Worlds. The steps taken to produce a prototype model using this type of package are described by describing them in terms of the creation of geometry, shading and lighting, the addition of texture mapping and the inclusion of behaviour.

5.6.1. Visualisation Geometry.

To build Virtual Worlds, developers choose between a range of modelling software packages that create and render high-quality images in three dimensions. These packages vary widely in ease of use and capability. During the course of this thesis, many VR modelling packages have started to emerge as commercial products. These products offer VE designers the opportunity to create, or import, worlds on commercially available personal computers. The early stages of the software analysis section was devoted to experimenting with as many VR modelling packages as possible to determine which packages would suit the aims of the project. The packages, which were utilised in these early experimentations, are briefly explained in Appendix C.

From a very early stage it became quite evident that many of the packages examined would not fulfill the needs of the project for several reasons. Some packages only provided crude representations with little or no texture mapping ability and some of those whose representations were of a high enough standard, did not provide images at a high enough frame rate to be useful. There were only two of the above programs that were considered of a high enough standard to merit further examination and they were the World Tool Kit for Windows from Sense8 and Virtus Walkthrough Version 1.0. It was decided to start work on Virtus Walkthrough V 1.0.

Of all the VR software that was tested, Virtus Walkthrough is the only product that produces instant VR. Most VR products assumes that a designer already knows something about VR, or that he is willing to spend a lot of time learning the software. Virtus Walkthrough uses a drag-and-drop approach to creating VEs. This makes it an ideal software package for an introduction to VR and VR techniques.

Several different versions of Virtus Walkthrough were used. The original version, Virtus Walkthrough Version 1.0., was the first version to be used. Then Virtus Walkthrough Pro 2.0. was used and, finally, a beta version of Virtus Walkthrough Pro. 2.6. was used. Although the original program is a very useful and versatile one, it has some very important limitations such as it does not allow for texture-mapping, collision detection or exporting 3D Data Exchange Format (DXF). Both the collision detection and the 3D DXF exporting facilities were enabled on Virtus Walkthrough Pro 2.0. but this did not allow for an adequate texture mapping facility when using Windows 95. This was rectified by the beta version of Virtus Walkthrough Pro. 2.6 along with some noticeable extra features such as its ability to export to the Virtual Reality Modelling Language (VRML) for use on the Internet.

Virtus Walkthrough is a computer aided visualisation package used to create 3D walkthroughs of any conceivable situation. Designing using this package is unlike conventional 3D modellers; it does not allow a designer to draw traditional lines and vectors. Instead, it takes an object modelling approach to draw 2D polygon-based objects with tools specifically developed for spatial design. The interface used in Virtus Walkthrough, illustrated in Figure 5.3., is divided up into three sections, a Design View, a Walk View and a Tools Window.

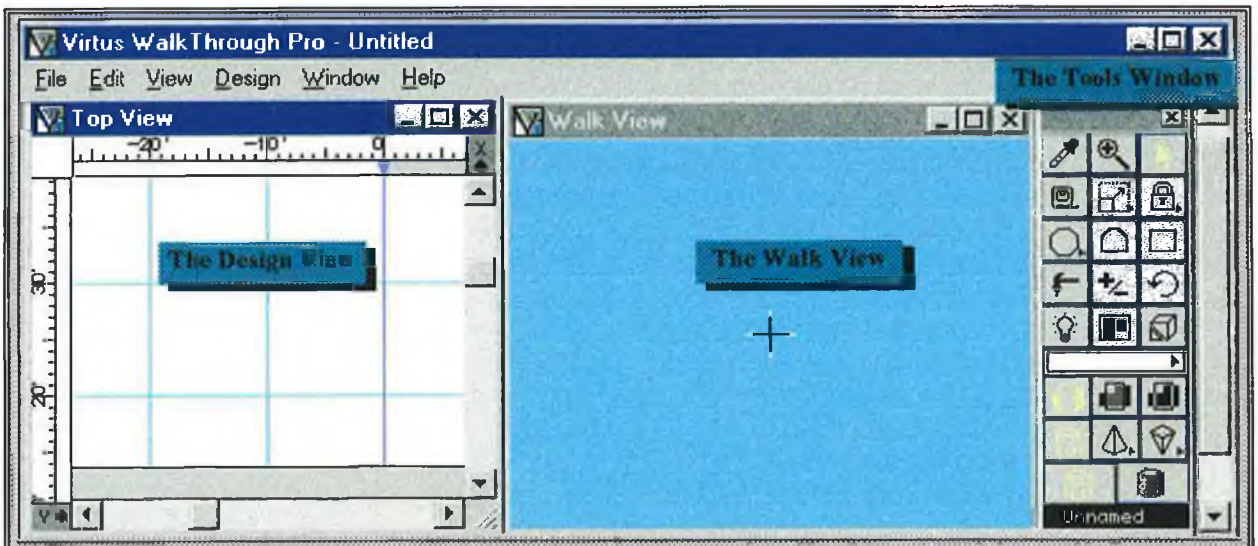

Figure 5.3. The Virtus Walkthrough Pro 2.6. Beta Interface.

The Design View is the 2D area where the designer lays down the basic dimensions of the objects within the Virtual World. There are six possible design views: top, bottom, left, right, front and back. A designer may draw, view or edit the drawings in any of the design views. If desired all six views may be opened at the one time. Multiple design views can be extremely useful and can help to increase the product's efficiency but in most cases, however, the size of a designer's monitor will limit the number of concurrent views. The Walk View enables a participant to walk through and around the 3D rendering of the objects that the designer creates in the design view. Finally, the Tools Window is the set of tools that a designer uses to create, and that a participant uses to navigate around, the environment. One of the rules in the Virtus Walkthrough manual states that "Speed and detail are inversely proportional" (Scott 1994). This is a very important rule and should be kept in mind at all times when designing using Virtus Walkthrough.

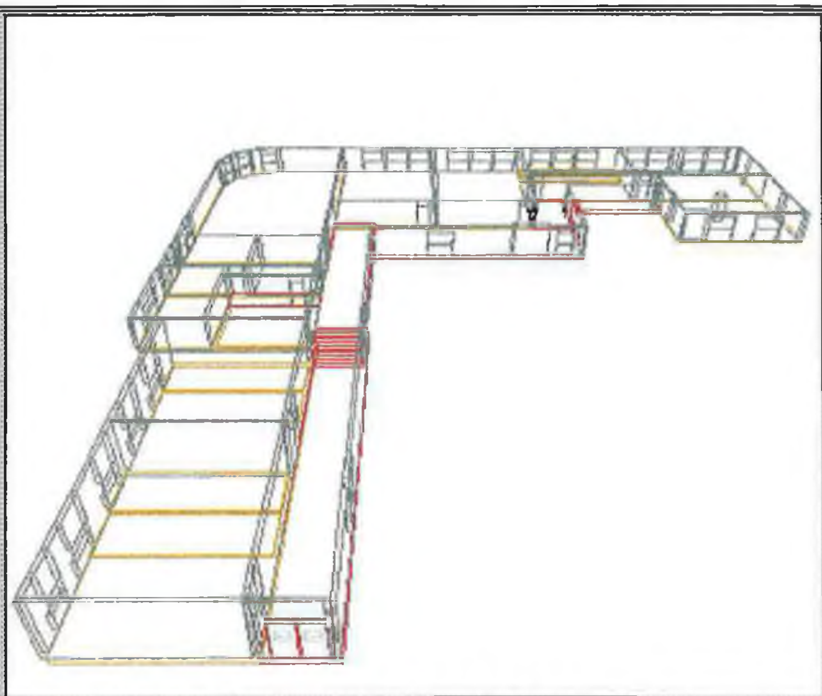
Figure 5.4. A Wireframe Model of the Computer Section, Cathal Brugha St.

Designing with Virtus Walkthrough was a relatively straightforward exercise. The first step was to design the basic structures of the Computer Section. This wireframe model, illustrated in Figure 5.4., was designed and stored in Virtus Walkthrough .WLK file format. The prototype model was not constructed to scale as this was not the main priority at this stage of development. Designing a model to scale with Virtus Walkthrough, however, was not a problem because it provides the designer with a ruler on the X, Y and Z axes which allows measurement in any units, thus, making scaling relatively straightforward.

Although Virtus Walkthrough has comprehensive furniture libraries it was decided that these libraries would not be suitable for inclusion into the prototype model as they were not the correct types of furniture and exact replicas of the real-life furniture were required in order to ascertain whether the construction of real life situations was feasible with Virtus Walkthrough. Therefore, every piece of furniture present in the Computer Section, from the computers to the classroom teaching tools, had to be designed. These pieces of furniture were designed as wireframe models, Figure 5.5., and were stored in a 3D library from which they could be entered into future models. This process was used to a significant degree in the construction of the prototype model as a large amount of repetition of furniture occurred in the Computer Section.

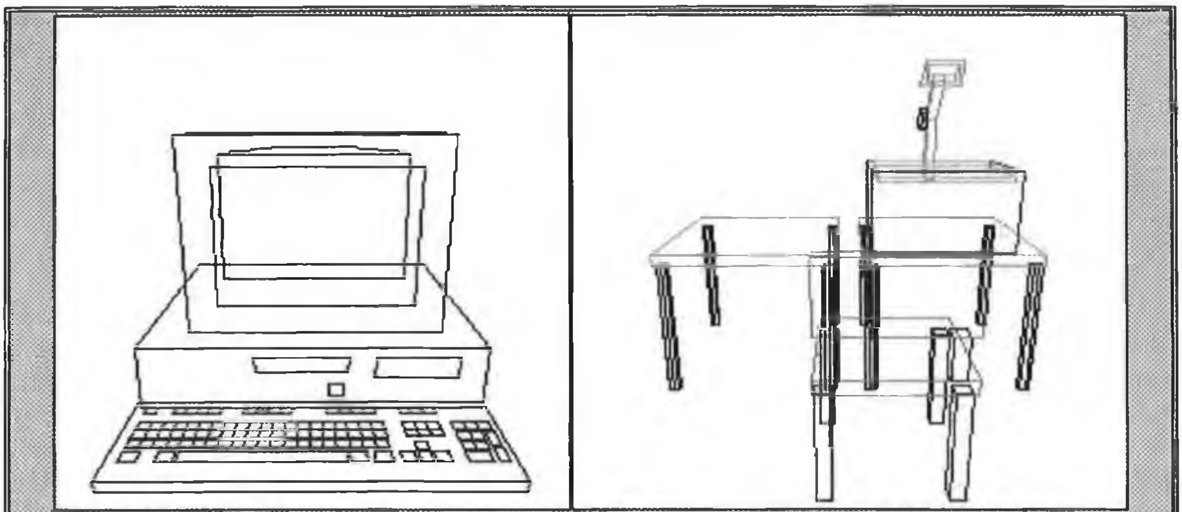

Figure 5.5. A Personal Computer and Teaching Tools Designed using Virtus Walkthrough 1.0.

At this stage in the development of the prototype model all of the basic wireframe objects had been constructed. This wireframe was used as a skeleton on which to build the worlds. The next step that took place in the geometry stage of the design process was to add the furniture, previously created and stored in the 3D libraries, to the model of the Computer Section, Figure 5.6. This was a very intricate and time-consuming process as furniture had to be placed very accurately into the prototype model.

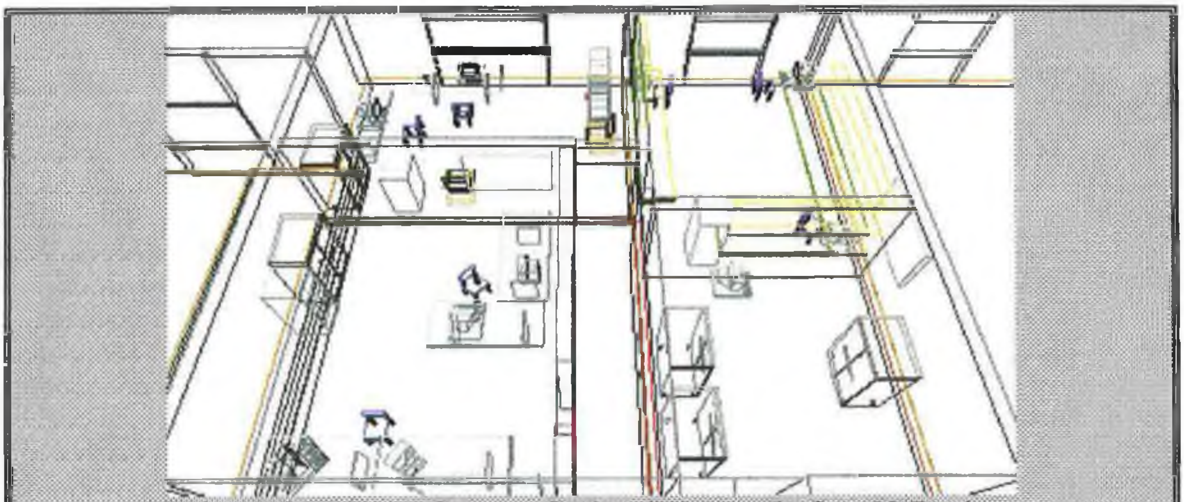

Figure 5.6. Wireframe of the Offices, Computer Section, Cathal Brugha St., with Furniture.

5.6.2. Visualisation Lighting and Shading.

Thus far, all of the graphical diagrams of Virtus Walkthrough have all been represented by wireframe models. In this section all of the basic geometry constructed as described in section 5.6.1. was enhanced with the addition of light and shade. Using Virtus Walkthrough, the application of the different types of shading is a relatively straightforward process. Firstly, simple white fill was added to the prototype model. This process fills the faces of the walls with a plain white colour and does not allow the inclusion of lighting effects. The view illustrated in Figure 5.7. provides a better reproduction of the building and the furniture within it. White fill is achieved by changing the shading and drawing options in the rendering properties of the Virtual World.

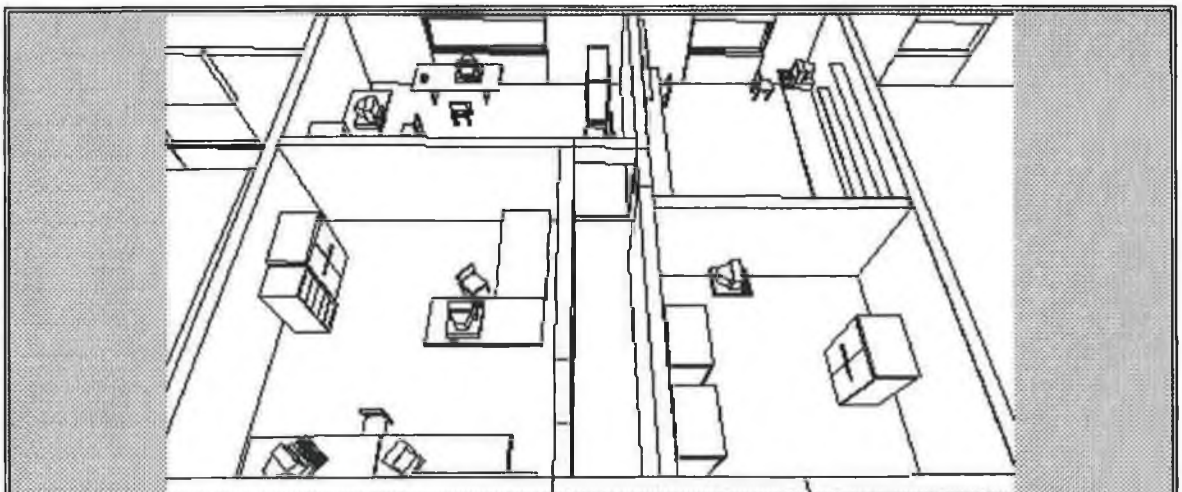

Figure 5.7. A White Fill of the Offices, Computer Section, Cathal Brugha St.

The second alteration to the geometry of the Computer Section was the addition of an unshaded fill, Figure 5.8. An unshaded fill allows a faces to display simple colours for the first time but still does not allow the inclusion of shading to these faces. The introduction of even simple colour to a VE provides a participant with a better sense of perspective of that environment. This type of fill is also achieved as described above.

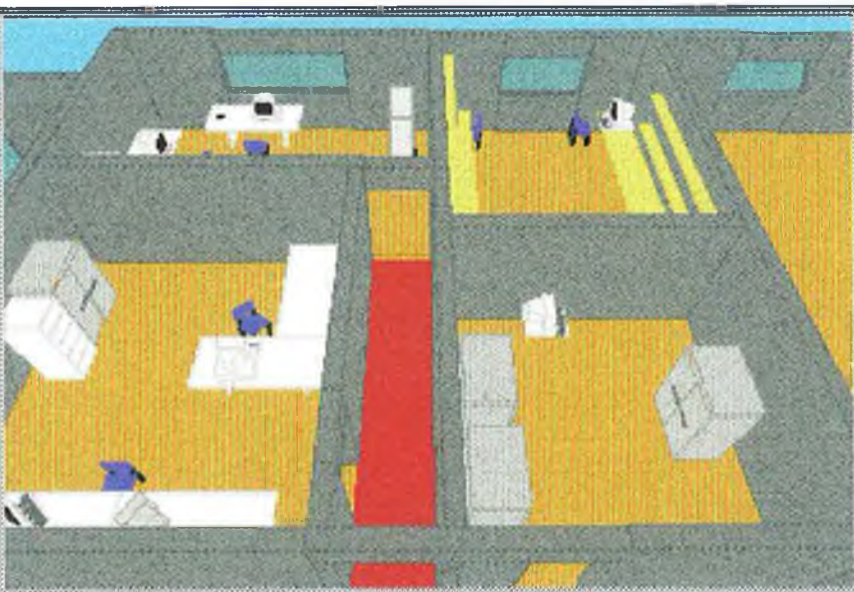

Figure 5.8. An Unshaded Fill of the Offices, Computer Section, Cathal Brugha St.

A shaded fill is the next type of shading to be applied. This type of fill adds a sense of depth and a greater sense of realism to the whole model. This type of shading, illustrated in Figure 5.9., allows a designer to display objects in a Virtual World with colour and allows the inclusion of shading into the model.

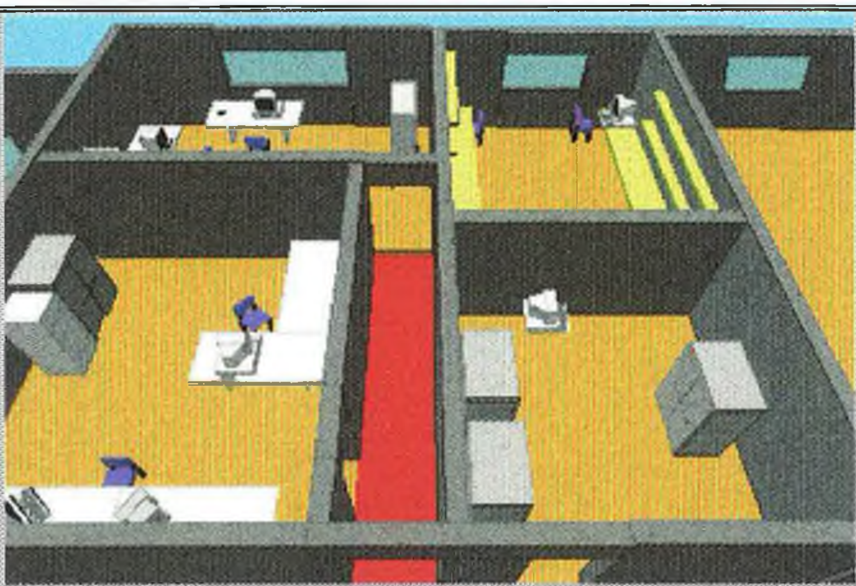

Figure 5.9. A Shaded Fill of the Offices with furniture, Computer Section, Cathal Brugha St

The next step was to add lighting effects to the prototype model. The lighting editor option in Virtus Walkthrough enables a designer to add lighting to a Virtual World. Lights in Virtus can be applied to an individual object or a group of objects, allowing complex combinations of lighting arrangements to be added to a world. Virtus Walkthrough allows for the addition of both ambient and directional light sources. As previously discussed in section 5.3.2. ambient light provided an equal amount of light to all parts of the world regardless of its position in the world. Therefore, only one ambient light is necessary. In Virtus Walkthrough the use of more than one light is possible but their colours and intensity will either cancel each other out causing white light or they will have an almost undetectable effect on one another. Directional light, on the other hand, has to be positioned in a world so that it shines on a specific object. The complexity of the lighting effects produced is directly proportional to the number of directional light sources present in the Virtual World. Both types of lighting effects have been applied to Figure 5.10., in an attempt to improve the atmosphere of the prototype model.

Figure 5.10. The Offices with Lighting Applied, Computer Section, Cathal Brugha St.

5.6.3. Visualisation Textures.

Actual images taken from the Computer Section should have been applied at this stage to the model in the form of texture maps. This did not occur in the case of the modelling and visualisation programs because at time that the prototype model was implemented the only version of Virtus Walkthrough that was being used was Version 1.0. which did not support texture-mapping.

5.6.4. Visualisation Behavioural Characteristics.

Virtus Walkthrough version 1.0. does not support the addition of any behavioural characteristics whatsoever. Though this changes slightly in Virtus Walkthrough Pro 2.0 and Virtus Walkthrough 2.6 beta these packages mainly devote their attentions to the appearance of a Virtual World. Virtus Walkthrough defines objects geometry, lights and textures, but lacks the important features of interactivity, such as object behaviours and cause-and-effect relationships.

5.7. Engines Used in Three Dimensional Games.

The second group of packages chosen to construct prototype models were not specific design programs at all: they were engines from well known 3D games. These engines were chosen because of their ability to allow the participant to interact within a VE in real-time. The steps taken to produce a prototype model using games engines are best described by describing them in terms of the creation of geometry, shading and lighting, the addition of texture mapping and the inclusion of behaviour.

5.7.1. Games Geometry.

The first step in this process was a comprehensive analysis of the different types of games engines available in an attempt to determine which, if any, of these engines would be appropriate for the construction of VEs for use in the tourism industry. At the time of this preliminary investigation there were only a few games engines available that could be used for such purposes. The engines that were examined in this primary investigation are briefly explained in Appendix D.

The outcome of these preliminary investigations showed that the engine from Doom, hereafter referred to as the Id engine, the engine from Hexen, hereafter referred to as the Raven engine, and the engine from Duke Nukem 3D, hereafter referred to as the 3D Realms engine, all warranted further investigation. In an attempt to determine which of these three games engines would be the most appropriate for the tourism industry it was decided to construct the prototype model using all three of these engines. This process was assisted by the fact that converters were available to convert from geometry created using the Id Engine to both the Raven and the 3D Realms engine. The construction process used to build the prototype model using each of these three games engines will be described, starting with the Id engine.

The Id engine is not in itself a VR editor but with the addition of certain editing tools the Id engine could become a very effective VR Tool. What makes the creation of VEs possible using the Id engine is that the creators of Doom, Id Software, released information about the game that made it possible for others to create editors and produce their own levels of Doom. Using these, a designer can not only create a totally new world but can also ascribe behaviour to each object within the world. The construction of a VE using the Id engine relies on creating vertices, lines, and finally, sectors.

Vertices - By definition, *"vertex"* means the termination or intersection of lines or curves. Vertices in the Id engine sense are points at which the lines converge. Certain editors allow a designer to enter individual vertices; other editors insert vertices with the sectors. Lining up vertices is relatively straight forward using most editors as they display grids which allow vertices to be aligned and the distance between them measured. Some grids are adjustable and scaleable, which enables the placing of vertices in exact locations. This can be a very useful facility when placing textures between two vertices without truncating the texture. The number of vertices that can be used in a world is limited by the specific editor used rather than by the Id engine itself. Vertices are the basic building block of a Virtual World with the Id engine.

Lines – *"Lines"* are marks on the map running between two vertices. A line simply cannot exist without a vertex at each end. Lines in the Id engine are where you assign textures for walls and behaviour for objects such as lights and elevators. Lines have a vast amount of information associated with them and it is, therefore, essential to have a good understanding of them in order to build a good quality Virtual World. Most editors use techniques that enable a designer to first place vertices in the world and then add a line between the two vertices but some editors insert sectors into the world which contain vertices and lines. Lines have many different attributes which could be assigned to them in order to achieve a different objective and could, therefore, be used in different situations.

Sectors - The next concept that needs to be understood about the Id engine is that of *"sectors"*. A sector is basically a group of lines and vertices that work together as a unit. Each sector has a unique number and attributes may be assigned to each to achieve a certain effect. Sectors also control the floor and ceiling height and the textures assigned to them. All objects in the Id engine are constructed of sectors.

Depending on the editor used, sectors will be constructed in different ways. Some editors enable a designer to insert a whole sector, while in other editors the designer must manually construct a sector by placing vertices and lines into the world and then assign them to a sector. Each sector has a specific set of attributes, which tells the Id engine how the space should appear and behave during a walkthrough. A sector's main attributes consist of the floor and ceiling details, lighting levels and behavioural characteristics. These three attributes are explained in more detail in Appendix E.

Sprites - A *"sprite"* is an object which can perform an action of its own or can interact with, or perform an action on, a participant. VEs created using the Id engine can also contain a number of items which are not created using vertices, lines and sectors. These items are called sprites. They are placed into an environment in order to indicate where they will appear when a participant enters it. Sprites include a wide variety of items such as participant's start positions, pieces of furniture, portal positions and computer operated characters. The only essential sprite that every world created in the Id engine must possess is a participant start position, otherwise, it is not possible to enter the environment. Each sprite possesses two attributes that can be set by the designer. These attributes are:

1. *The Facing Angle* attribute determines the direction in which the sprite will be facing when a participant enters the world. Only participants' start positions, computer operated characters start positions and portal destinations make use of this attribute; furniture sprites look the same from any angle.

2. *The Multi-Participant* attribute enables the use of sprites that will only appear in a multi-participant environment. This is useful for the addition of obstacles to limit participants' access to items such as portals to other environments.

There is a large number and variety of tools available to create new VEs, or alter existing ones, using the Id engine. Before deciding on the most appropriate tools to use all the tools should be examined and the best tool, or a mixture of tools, chosen. In an attempt to gain an understanding of the basic concepts of the Id engine and of the numerous editors that are available for it, it was decided to conduct preliminary experimentation on levels of Doom using a number of these editors. This preliminary experimentation is described in Appendix F.

As a result of these undertakings a package called Doom Editor Utility Version 5.1. (DEU) was chosen as it was found to be the most effective map editor available. DEU was to be used in conjunction with other tools such as the graphics editing utilities, Wacker and DMGraph, and the sound editing utilities, DMaud and DMMusic.

Using the Id engine and DEU, the design of the structure of the prototype model, demonstrated in Figure 5.11., was relatively straightforward. The only extra item in this model was that of the participant start positions as they are essential if the world is to be viewed with the Id engine. This model did not contain windows, doors, steps or different ceiling heights and in order to include them the model firstly had to be divided into 56 separate sectors. Each sector had to be allocated a specific floor and ceiling height, thus creating the basis for steps, windows, doorways, and ceiling heights. This complicated many things and extra time was spent adding Upper and Lower Textures to accommodate the different ceiling and floor heights. The outer wall of the room, in the upper right corner of Figures 5.11. and Figure 5.12., also needed adjustment. The next step that was taken was to introduce doors into the model. Initial problems arose but with several alterations to line and sector attributes these problems were solved and the structure appeared as it is in Figure 5.12.

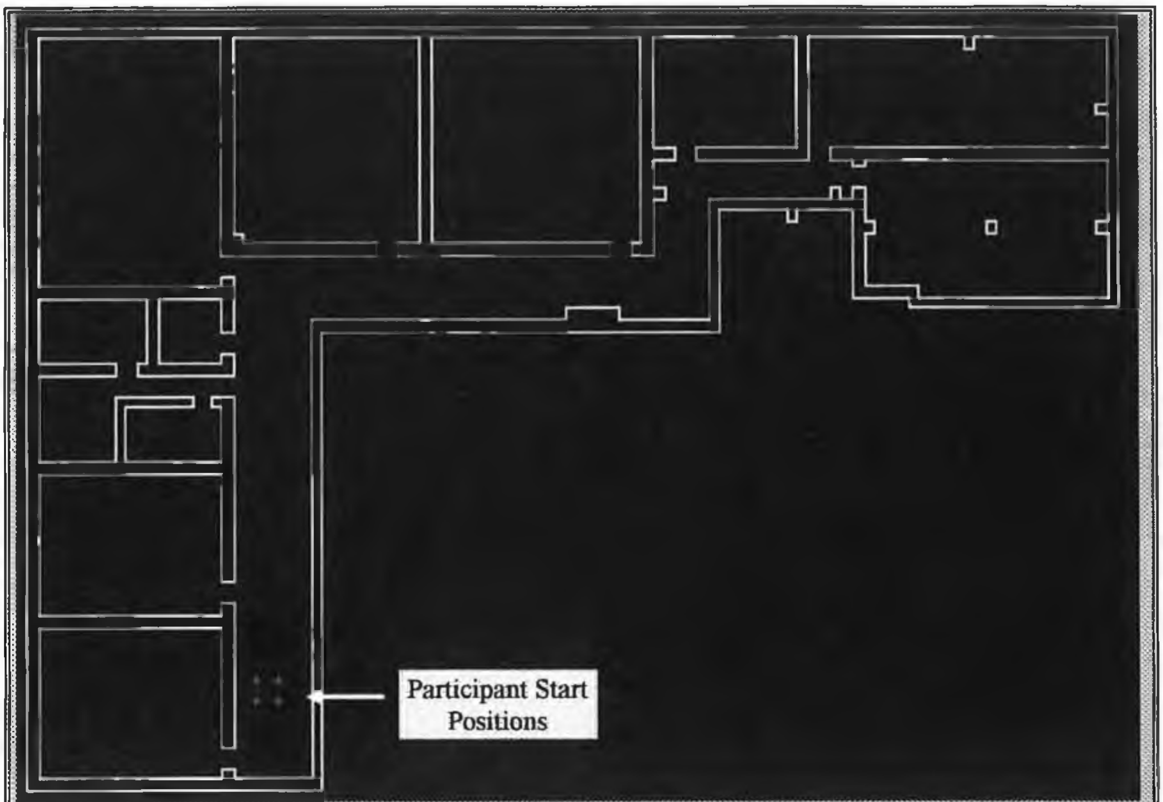

Figure 5.11. A 2D Map of the Prototype Model.

Internal objects such as furniture were now added. Problems arose here because certain pieces of furniture such as benches, tables and chairs required one sector to be placed above another and the Id engine would not permit this. A common misconception with the Id engine is that the Virtual Worlds produced are fully 3D. As a designer, it is essential to understand that the Id engine is what is normally termed pseudo-3D. The main consequence of this pseudo-3D nature of the Id engine is its inability to permit any location in the worlds to exist above another location. What this means from a designer's perspective is that no sector, or part of a sector, may overlap any other sector. This has huge practical implications in that one room can not be built on top of another room, bridges can not be built and roofs can not be built on buildings that a participant can enter. This makes many real world environments immediately un-designable.

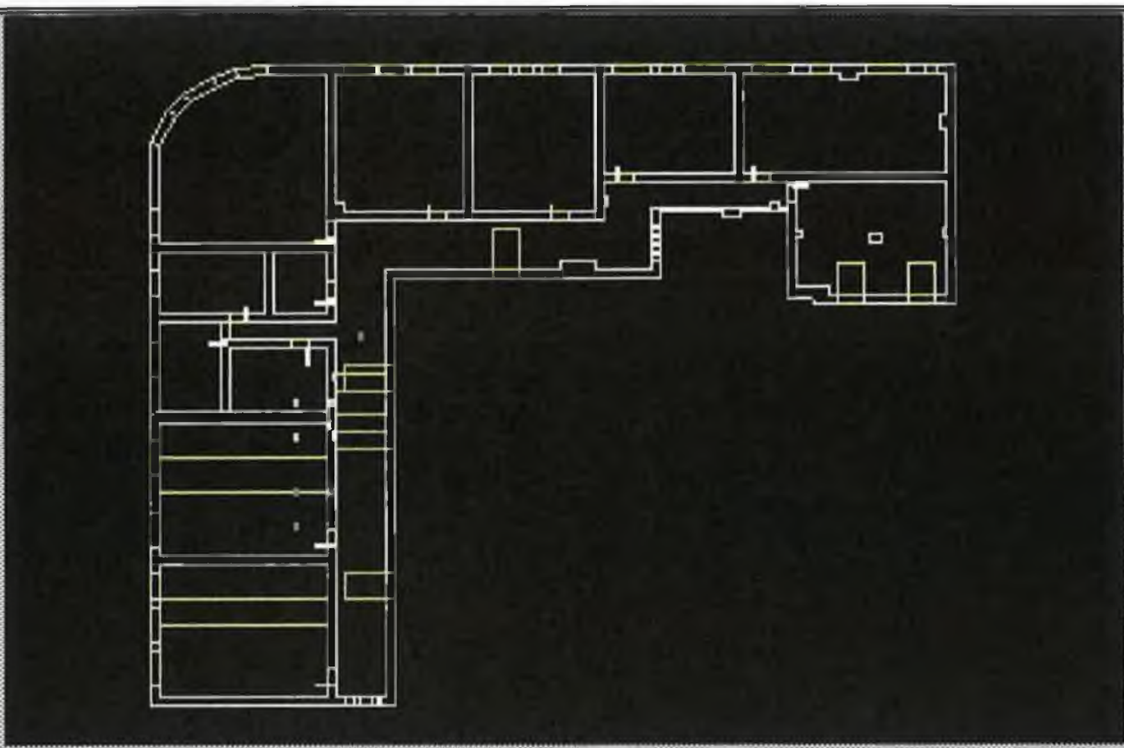

Figure 5.12. A 2D Map of the Prototype Model with Sectors Included.

In an attempt to combat this limitation most of the furniture required can be inserted into the structure in the form of sprites. This process requires a piece of furniture such as a chair or printer to be drawn with the use of a drawing package such as Photo-Paint to draw a 2D image of the object. The image must be saved to a .GIF or a .PPM file format for inclusion into the Id Engine. To input the image file into the Id engine a graphic program called DMGraph was used. Once an image had been entered into the Id engine it could then be placed into a world in the form of a sprite. The same procedure was used to create and place computer operated characters into a Virtual World. The character sprites are constructed using a series of images rather than just one. In the case of the character featured in Figure 5.13. it was constructed of 52 separate images of the character in varying positions, because that number of separate images is required in order to create a character which walks with a fluid movement.

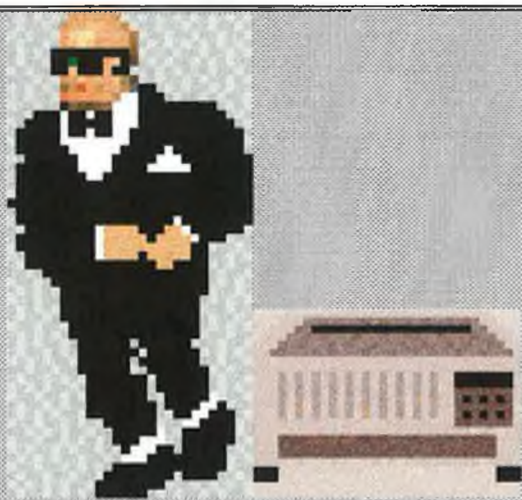

Figure 5.13. A Character Sprite (1 Image of 52) and a Printer Sprite

The furniture was then placed into the basic shell of the college by using vertices, lines and sectors where possible and sprites for other pieces of furniture that could not be modelled by those means. At this stage all the objects in the prototype model had been constructed using a wireframe representation, Figure 5.14. At this stage in the process the world could be converted from the Id Engine file format to the Raven and 3D Realms Engine file formats using the appropriate converters. Once the model had been converted to the appropriate engines it could then be viewed using these engines. The structure of all three models remained the same with regards to the vertices, lines and sectors used but the sprites that had been added to the Id engine could not be converted automatically to the other two engines. This process had to be conducted manually.

Figure 5.14 A Wireframe Representation of the Prototype Model (Basic Structure) using the Id Engine.

In the case of the Raven Engine the appropriate sprites were imported to the engine with the use of an editing utility called Wintex. Wintex is a utility which can be used with both the Id engine and the Raven engine. Once the appropriate sprites had been imported successfully they were placed into the prototype model. This was achieved by the use of a Map editing utility called "Hexen Editor for Total Headcases" (HETH). The HETH editor is, in fact, quite similar to DEU, and this drastically reduced the time spent learning the utility.

In order to import the appropriate sprites into the 3D Realms engine, a utility called Editart was used. Editart, as the name suggests, is a utility for editing graphics in the 3D Realms engine. After the sprites were imported into the engine they then had to be placed into the prototype model. This was achieved with the use of a map editing utility called Build. Build, unlike the map editors used for both the Id engine and the Raven engine, includes a 3D mode for editing the environment. In this editor a designer can place a object, not only in respect to its x and y coordinates, but also in respect to its z coordinate.

An object can, therefore, be placed precisely where the designer wishes to place it. Another advantage of this method of modelling is that a designer can walkthrough the world without having to exit from the editor and run the actual engine. This is both an extremely effective and efficient method of modelling. Once the furniture sprites were placed into the model their dimensions were then altered. They were stretched or squashed both horizontally and vertically in order to make them the correct size in comparison to their surroundings. Sprites in the 3D Realms engine have many additional features compared to their counterparts in the other two engines. They can be resized, shaded, lit up, placed on walls, in a horizontal position or permanently facing the one direction. These additional features have a large beneficial effect on the potential of the 3D Realms engine as a VR engine. These features enable a designer to build tables, bridges, rooms over rooms and many other things that are impossible with the other two engines. The same pieces of furniture were placed into the model as were in the two previous engines but this engine also allows a designer to add glass, tables (which a participant can now crawl under) and pieces of furniture which can be placed on walls such as fire extinguishers and clocks.

5.7.2. Games Shading and Lighting.

The next stage in the construction of prototype model using these games engines was to add shading to the wireframe representation. This process is facilitated by the fact that all three of these games engines have default textures. These textures are displayed in Figures 5.15.,5.16., and 5.17. Even though the entire Computer Section has been modelled with all three of these engines, in an attempt to facilitate the demonstration of the steps taken only one room, known as "Room 42", will be displayed from this stage on. Room 42, an office in the Computer Section, was chosen to demonstrate this because most of the objects present in the Computer Section are also present in this one small room.

Figure 5.15. The Computer Section, without Furniture, with the Id Engine's Default Shading.

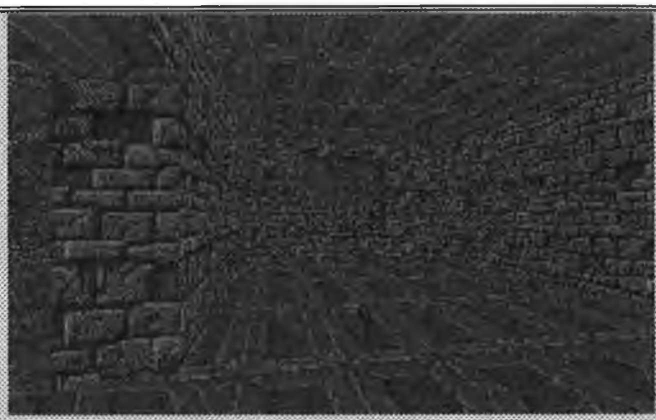

Figure 5.16. The Computer Section, without Furniture, with the Raven Engine's Default Shading.

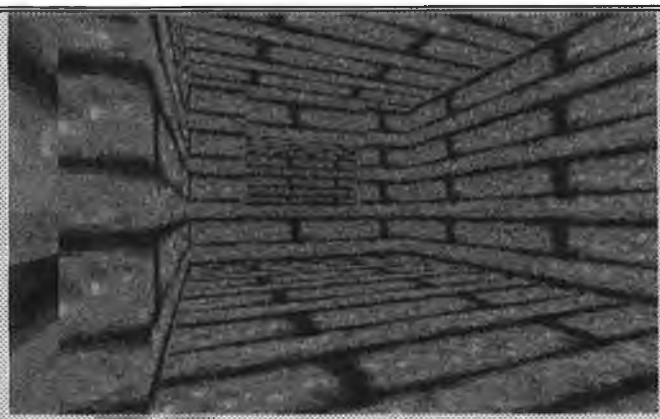

Figure 5.17. The Computer Section, without Furniture, with the 3D Realms Engine's Default Shading.

All three of the games engines take full advantage of lighting and lighting effects. Lights can vary from worlds with absolutely no light to worlds with very bright lights. The Id engine and the Raven engine both have 34 levels of lighting, ranging from a value of 0, which is complete darkness, to a value of 255, which is extreme brightness. Lighting levels increase incrementally every 8 units between the 0 to 255 range but values 100 and below are very dark and appear about the same. Lighting levels are set by attributing a certain value to a sector and all objects that appear within a sector would, therefore, be illuminated to that level.

Lighting levels in the 3D Realms engine is dealt with slightly differently. Instead of applying lighting to a whole sector, as is the case in the previous two engines, the 3D Realms engine applies light to each face of an object. This allows far more flexibility and a wider array of lighting combinations. This also allows a designer to experiment with shadows and more complex shading effects.

Designers can also set lights to blink, shimmer, pulsate or set them to be turned on or off at the flick of a switch. They can even simulate night and day and the effect of light diminishing as one moves further away from the light source in the Virtual World. These special lighting effects can be achieved by all three engines although they do vary in the techniques that are used to achieve them. In the Id engine and the Raven engine lighting effects are set by assigning certain attributes to a sector whereas in the 3D Realms engine these effects are set by changing the characteristics of a sprite placed within a sector. All of these light sources are generated with reference to the brightness level of the sector and the brightness level of adjacent sectors. The value specified for the sector's own brightness level determines the upper limit of the fluctuation, while its lowest level is taken from the lowest light level amongst all immediately adjacent sectors. If there are no adjacent sectors with a lower lighting setting the light will usually fluctuate between its light setting and total darkness. In the case of the prototype model, light switches, illustrated in Figure 5.18., were placed on the corridor and in all of the rooms. These switches worked in the exact same manner as they would in a real-life situation. The switches in the Id engine and the Raven engine were entered as wall textures and the light switch in the 3D Realms engine was a sprite applied to an appropriate position on a wall. In the case of the technique used in the Id and Raven engines the texture had to be aligned carefully and often took several applications to perfect. The latter technique was far more straightforward as the sprite once applied could then be manipulated easily.

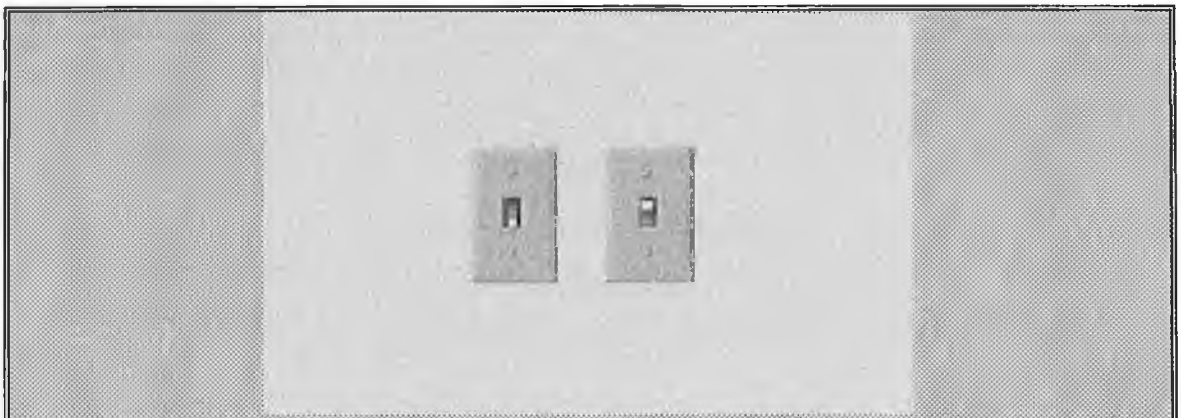

Figure 5.18. Two Light Switch Textures, First One On and Second Off, Used in the Prototype Model

5.7.3. Games Textures.

The next step involved choosing the correct wall, floor, and ceiling textures. The entire section dealing with creating and importing textures into the games engines proved to be a time-consuming exercise. This was mainly due to the fact that there was very little literature on the subject and a trial-and-error approach had to be used to solve many of the problems that arose. The main aim here was to introduce appropriate textures into the Virtual Worlds created in an attempt to make them more realistic to the participant.

At this stage in the process textures taken from the real walls in the Computer Section were added. The textures were obtained using two methods - video and photographs. The aim was to make the photographs and the video clear from obstruction, shadow free, square to the viewer and with as little distortion as possible. The video was recorded using a Panasonic NV-M40 VHS Movie Camera was edited on a Panasonic VCR - AG5700 and was captured to the computer using a video capturing card and a copy of Asymetrix Digital Video Producer (DVP) Version 3.0. The photographs were taken using a Pentax camera, were then developed and magnified using an OCE 3107C colour photocopier. This helped to get a better wall texture and helped to reduce any shine on the photographs. These magnified photocopies were then scanned at 600 Dots Per Inch (DPI) with a HP Scanjet 4C scanner. A software package called Corel Photo-Paint Version 6.0. was used to straighten the images and crop off what was not needed. These images were then stored as Windows Bitmap files and then converted to the correct format for inclusion into the games engines. The reason that both methods were used was in an attempt to ascertain which method would be of the best quality for the use with the models. The photography method produced a far higher quality texture and it was, therefore, decided to use this method to produce future textures.

Once the photographs had been stored in a suitable format on the computer. The next step was to ascertain whether textures could be imported into the three engines and if they could, what format they needed to be in order to do so. In order to describe this process it is best deal with each engine individually.

5.7.3.1. Textures - The Id Engine.

The first experimentation that was conducted in this area was carried out using a beta texture editor called Wacker. Textures were created using Wacker and were placed into the Id engine. These textures were subsequently applied to appropriate objects using DEU. Everything seemed quite straightforward initially but as soon as the Virtual World was started either the engine crashed or the new texture caused an undesirable effect. Several other attempts were made to create textures using Wacker but they were all to no avail. These attempts included changing existing textures, importing plain colours and, finally, importing graphics files in different resolutions and formats such as Windows bitmaps, .BMP, and graphics information files, .GIF.

The next approach used a graphics editor called DMGraph. DMGraph enables a designer to edit existing textures in the Id engine or create and enter totally new textures into the engine. The first experiment carried out using DMGraph was where textures were exported from the Id engine in .GIF format and were reintroduced into the Id Engine using a different name. The purpose of conducting this seemingly futile exercise was to examine the steps required to import and export textures using this editor. These initial exercises were successful, thus, proving that such a method could be used to introduce a new texture into the Id Engine.

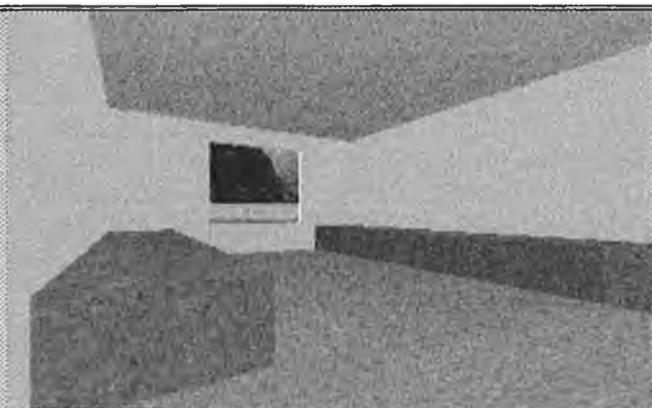

Figure 5.19. An Example of Textures Mapped onto the Prototype Model in the Id Engine.

The next step was to attempt to introduce a totally new texture into the Id engine. Textures imported into the engine must adhere to certain requirements. The file can only be a .GIF or .PPM format file with a resolution of 320 x 200 or less and a maximum of 256 colours. In order to add a partially invisible texture, the invisible section has got to be set to the colour cyan (set the red value to zero, maximise the green value and set the blue value to 255). An example of actual textures applied to the prototype model is displayed in Figure 5.19.

5.7.3.2. Textures - The Raven Engine.

This section deals with the steps taken in an attempt to introduce textures into the Raven engine. Wintex was used to introduce textures into the engine. The Raven engine currently accepts only 320 x 200 resolution, 256 colours Windows Bitmap pictures. Wintex will convert 24 bit (millions of colours) bitmaps, to the appropriate 8 bit bitmap but the conversion will be extremely slow and, because it does not support dithering, may be unsatisfactory. Corel Photo-Paint was, therefore, used for bitmap conversion. The textures were then added to the prototype model in the Raven engine using the HETH editor. An example of the prototype model, complete with the actual textures, is shown in Figure 5.20.

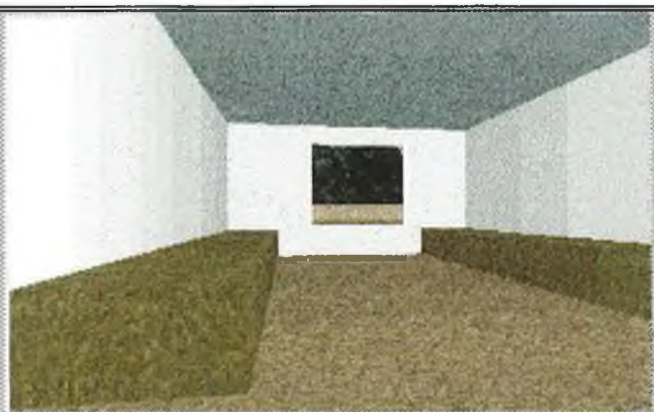

Figure 5.20. An Example of Textures Mapped onto the Prototype Model in the Raven Engine.

5.7.3.3. Textures - The 3D Realms Engine.

In this section the introduction of textures into the 3D Realms engine is described. The main tools used in this process were a texture-editing tool called Editart and a map-editing tool called build. Both of these tools are extremely reliable and well documented. This simplified the process significantly. Editart will accept 320 x 200 resolution, 256 colour Windows bitmaps or .GIF format files. Once the texture files had been imported into the 3D Realms engine it was then time to apply them to the prototype model, Figure 5.21.

Figure 5.21. An Example of Textures Mapped onto the Prototype Model in the 3D Realms Engine.

5.7.4. Games Behavioural Characteristics.

This section describes the concepts concerned with the application of behavioural characteristics to the VEs constructed using the Id, Raven and 3D Realms engines. This is the area in which these three games engines vary most and for that reason each engine is dealt with separately.

5.7.4.1. Behaviour - The Id Engine.

For behaviour, the Id engine utilises line and sector specials, which are fixed actions assigned to lines or sectors. These attributes enable the designer to set several different behavioural situations such as elevators, portals and a host of lighting effects.

Lines and sectors provide the Id engine with vital information about the disposition and appearance of the edge of sectors. Lines, however, can play a secondary role, they can act as triggers. Each line, in addition to its normal appearance attributes, can also carry a special action attribute. This attribute determines the nature of the action the line will perform once the participant crosses this line. These actions are set as code numbers which each perform a specific action. Usually a line will have a number set to 0, which means that no special action will be carried out by this line. The actions triggered by lines are determined by the Id engine. A designer can only utilise the actions provided and these actions can not be altered or added to. Lines cannot initiate more than one special action, although they may initiate that action in more than one location.

It is possible to use lines specifically for their trigger action and, therefore, not allow them to contribute to the provision of sector information. In this way trigger lines can be laid out in any position across a sector. The usual rules for lines still apply; they must run from vertex to vertex and they cannot cross other lines.

The actions triggered by lines are generally performed on some associated sector. The process by which the Id engine selects which sector to affect varies from action to action but generally falls into two basic categories:

1. *Local action mechanisms* operate on the sector that owns the left side of the active line, and is obvious when placing a line. These actions all provide doors to open. If this were the only activation mechanism available then the Id engine would only provide limited behavioural attributes.

2. Remote action mechanisms use tags to enable a designer to produce more complicated and useful behaviour in a VE. This is an additional way that the Id engine uses to connect special lines to the sectors they affect. The creation of remote actions is quite simple. In addition to having a special attribute, each line carries a tag number. This tag number is nothing more than an arbitrary identification number. Ordinarily, this tag number will be set to 0, meaning that the tag is not in use. When a designer wishes to create a remote action he must first of all set this tag number to any number other than 0 and he must also change the tag number of the sector, or sectors, that he wishes

this action to have an affect on, to that number. Actions that utilise a tagging mechanism operate by matching the tag of the activating line to that of a target sector, or sectors, whenever the line is triggered. Therefore, a single line can initiate its action in several sectors simultaneously, simply by having the sectors share the same tag number.

In addition to having two types of activation mechanisms, there are many different types of action which can be performed by special lines. These actions all follow the same basic principles for construction. The way that these actions are activated depends on their activation type. There are four activation types - manual, spacebar, walkthrough and impact. These activation types are explained fully in Appendix G.

Additionally, each special action may be repeatable or may be once off occurrences depending on the line's special attributes. In the case of the prototype model light switches were included in each of the rooms in the Computer Section to allow for the lights in each room to be switch on or of depending on their current status. This was achieved by using spacebar activation and a special texture, Figure 5.18., was imported to deal with the actual light switch. A repeatable trigger was used to allow the lights to be turned on and off several times. Certain voice-overs were placed into the prototype model to give the participant additional information about each room as he entered it. The voice-overs recorded in Creative Wave Studios were added to the Id engine with the use of a sound editing utility named DMaud. The Sound files were placed into the models and could be activated by several different means. This was conducted in order to first of all learn the different types of activation mechanisms and secondly, to ascertain which, if any, would be the best type to use in future tourism based models.

5.7.4.2. Behaviour- The Raven Engine.

The Raven Engine uses the same principles of activation mechanisms and types of action as the Id engine. The main difference between the two engines with regard to behaviour is that in the Raven engine instead of the actions being determined be the engine alone, as they are in the Id engine, a scripting language is used to enable a designer to modify existing action or create completely new ones. This scripting language is called Action Code Script (ACS). Each world has an ACS file that contains the scripts specific to that world. The scripts within it are identified using numbers that the general special ACS_Execute() uses. A script itself can call the ACS_Execute(), which will spawn another script that will run concurrently with the rest of the scripts. A script can also be declared as OPEN, which will make it run automatically upon entering the Virtual World. A script can be written in plain text. This script must then be saved with a .ACS file extension. This file is then compiled to a .O extension by using a DOS based program ACC.exe. This program will produce an output file, 'filename.o', from the input file, 'filename.acs'. The contents of the output file can then be inserted into the Behaviour lump of the worlds that it is used with. In order to do this the output file first has to be converted from a .O file extension to a .RAW file extension. Then the scripts can be inserted into the Behaviour lump of the Virtual World using HETH. The scripts, once inserted, can be executed in the exact same manner as they are in the Id engine, or by using an OPEN script which will be executed as soon as a participant enters an environment. The fact that the Raven engine can use a scripting language to assign new actions to the Virtual Worlds is a major advantage because it enables the designer to supply the participant with more information about the environment at different places during his experience.

In the case of the prototype model, the Raven engine was also used to set up light switches in each of the rooms in the Computer Section, using the exact same mechanisms as were used in the Id engine. The voice-overs were entered into the Raven Engine by using Wintex and then they were placed into the appropriate positions throughout the model by using the map editor HETH. The only difference that was made to the prototype model with the Raven engine was that an OPEN script was used to play a sound file as the participant entered the model.

5.7.4.3. Behaviour - The 3D Realms Engine.

The application of behaviour to a Virtual World constructed using the 3D Realms again uses much the same principles as the Id engine. Although the 3D Realms engine does not allow a designer to enter customised actions scripts into it, the engine does provide a far wider range of actions to deal with from the start. The activation mechanism used by the 3D Realms engine is a sprite used as a sector effector. Sector effectors define what action will occur in the sector that they are placed in when the action is summoned. These sector effectors are called upon by changing the Hi-tag and Lo-tag attributes of the sector and the sprite. Lo-tags define which action will be performed in much the same way as special action attributes do in the Id engine. Hi-tags, on the other hand, define where exactly this action is going to occur in the same way that tags do in the Id engine. Depending on the action that is to be performed some effectors need both a Hi-tag and a Lo-tag, while others only need a Lo-tag.

The same behavioural characteristics were added to the model as were applied with the other two engines. The light switches were applied in much the same manner except that the switches in this engine were applied as sprites as opposed to being applied as wall textures in the other two engines. Sound files were also included except in this engine the files, again placed as sprites, could be located in a specific location. In additional to these behavioural attributes many others were added in the case of the 3D Realms engine. Such features as swinging doors on each room, working video cameras, breakable glass in all the windows, clocks that break on impact and bins that dent when a participant punches it. All these features combine to bring a greater sense of realism to the environment.

5.8. Three-Dimensional Graphics and Animation Packages.

The third category of packages chosen was high quality three-dimensional graphics and animation packages. This type of package provides the user with a highly vivid, pre-programmed computer-generated walkthrough. The main drawback of such packages is that they do not allow the user to interact with the VE. It is, therefore, a thoroughly pre-programmed experience. Strictly speaking, of course, animation is not VR at all but it is being dealt with at the same time as the other VR techniques because 3D animations of today is where the ideas for VR of the future will be conceived, designed and developed (Wodaski 1996). The road to creating VR is a long one and most of the key steps along the way will happen first in the world of 3D modelling and animation.

Until relatively recently, the cost of first-class 3D animation has been prohibitively high (Wodaski 1996). At the high end of the animation spectrum are the animations like those produced for the films "*Jurassic Park*", "*Terminator II*", "*Lawn Mower Man*", and "*Johnny Mnemonic*". At the low end are animations that are often used in commercials. In general the more one spends, the more realistic and flexible the animations become. This scenario, however, is rapidly changing. Although 3D Studio, the package that is being used to create the animations in D-VR3, would be considered expensive by normal software standards, it is considered extremely inexpensive by 3D animation standards. With each release 3D Studio has moved desktop animation closer and closer to the high-end animation creation.

5.8.1. Animation Geometry.

The first step in constructing the prototype model is to use the standard 3D Studio tools to create the various objects that make up the component parts of the Computer Section. To create an object using 3D Studio, a designer would select from a large number of options to create different types of geometry. Once the designer selects the type of geometry he wishes to use, he then has to define where he wishes to place the object in the environment. He does so by assigning coordinates to the object. All objects are created using the same principle and once they are created they can be manipulated in a number of ways to display the type of object the designer wishes. A wireframe representation of Room 41, an all purpose class room in the Computer Section, which was constructed using 3D Studio is illustrated in Figure 5.22.

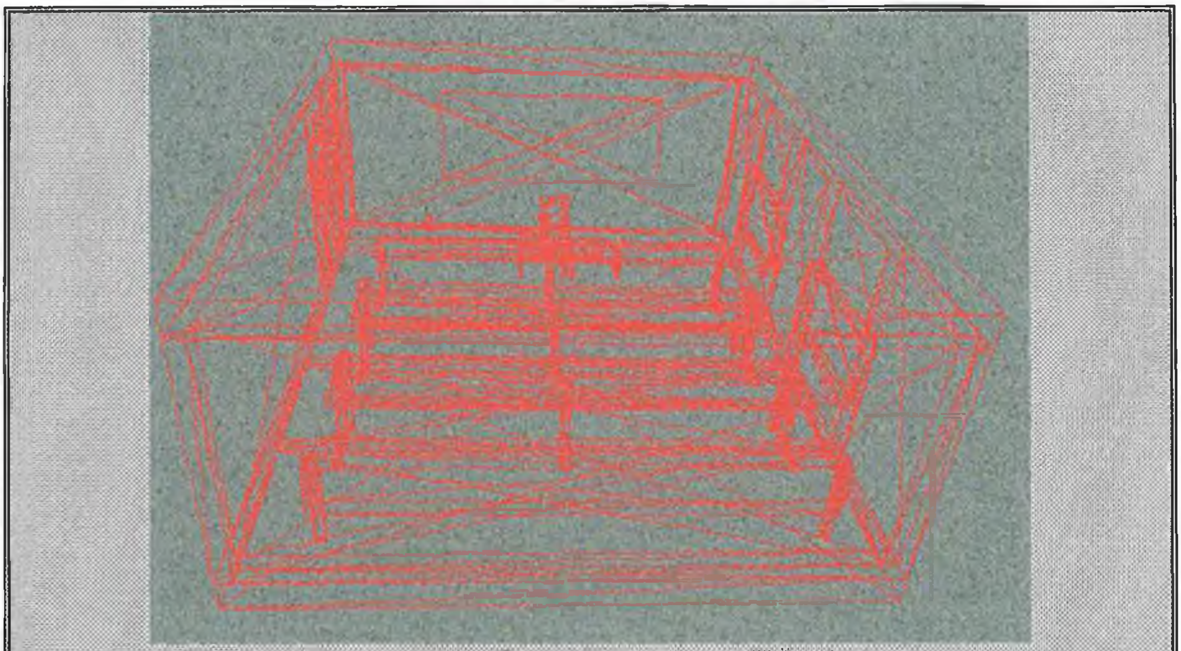

Figure 5.22. A Wireframe Representation of Room 41 Constructed Using 3D Studio

With this process a designer can create a complete VE regardless of the complexity but there are easier ways of achieving this goal. A designer can import .DXF files created with other software packages. In the case of D-VR3 the environments can be created in Virtus Walkthrough and exported to 3D Studio using a 3D DXF file format. This process saves all the repetition of creating the file in 3D Studio. The DXF files are the standard of the 3D universe. These files are commonly used by AutoCAD but are also supported by a wide variety of software packages.

5.8.2. Animation Shading and Lighting.

Although the structure of the environment is present, there are no surface shades, lights, camera viewpoints or textures on any of the objects. 3D Studio allows a designer to associate specific materials with objects. In fact it has its very own material editor which can be used to specify the characteristics of a wide variety of materials.

By default the ambient light in a scene is set to 77. There are 256 possible values, from 0 to 255: a value of 77 represents 30 percent of full lighting. The default value is displayed if the Lights/Ambient from the options menu is chosen. For an adequate light level a range of between 200 and 255 is advisable. The colour of ambient light can also be changed.

3D Studio also enables a designer to create two special types of lights: omni-lights, which are point sources of light; and spot-lights, which cast shadows, have adjustable light cones and allow the light source to be used as a projector to play video files or animation files from the source. By using several different kinds of lighting in different positions, intricate lighting arrangements can be achieved. Figure 5.23. illustrates Room 41 with appropriate shading and lighting.

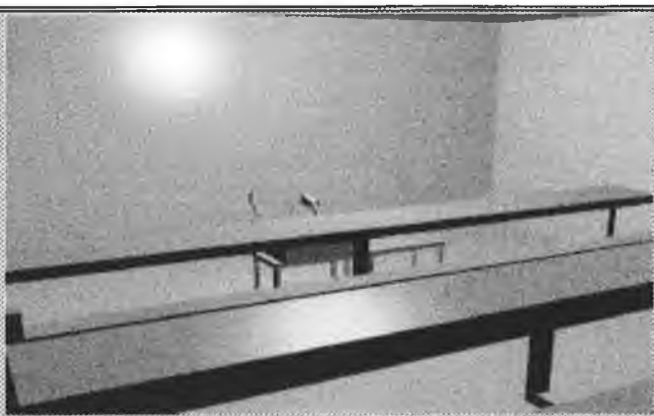

Figure 5.23. Room 41 with Shading and Lighting Constructed Using 3D Studio.

5.8.3. Animation Textures.

To apply a texture map a designer needs to define mapping coordinates so that 3D Studio knows exactly where to position the texture map on the objects and at what scale. The easiest way to do this is to fit the size of the map to the entire object by using the Surface/Mapping/Aspect/Region Fit option to display the mapping coordinates. However, the easiest way does not always provide the desired results. Therefore, other options, such as Bitmap Fit or Object Height and Width, might have to be used. Once the mapping coordinates have been defined the next step to take is to apply them to the specific objects by selecting the object and then using the Surface/Mapping/Apply/Object menu option.

After the mapping coordinates have been added, the next step is to apply a material. This material, in effect, will paste the texture onto the surface of an object. Creating custom material is a relatively straightforward exercise with the material editor in 3D Studio. Any existing texture, video clip or animation can be placed onto the surface of an object. The material editor also allows a designer to change the characteristics of a texture such as shininess, rendering quality, transparency and bumpiness.

Once a material has been created, applying it to an object is relatively simple. With the Surface/Material/Choose option a designer can choose the texture he wishes to use from a list and then with the Surface/Material/Assign/Object menu option apply it to a specific object or selection of objects. An example Room 41, complete with textures and directional lighting, is illustrated in Figure 5.24.

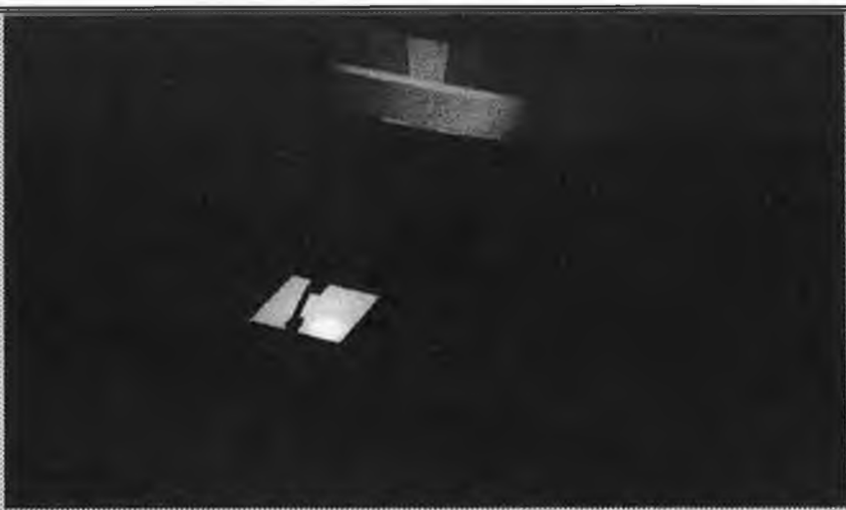

Figure 5.24. Room 41 with Textures and Directional Lighting Constructed using 3D Studio.

5.8.4. Animation Behavioural Characteristics.

VEs created in 3D Studio are non-interactive and the only behaviour that they possess is through the creation of animation. The animation process in 3D Studio takes place in the Keyframer. A designer does not have to create each frame of the animation when using the Keyframer. In fact, as the name suggests, only certain key frames in the animation have to be created and the Keyframer will generate all the intervening frames.

Table 5.1. Rendering Palette Choices.

Choice	Description
Low	The renderer creates an optimal 256 colour palette for the first frame and then uses that palette for all subsequent frames. Unless the first frame is highly representative of the remaining frames this is not a good option.
Medium	The renderer creates an optimal 256-colour palette for each frame in the animation and then combines the palettes into one optimal palette.
High	Each Frame is rendered with 24-bit colour. After rendering is complete, a single, 256-colour palette is calculated. This option takes the most time to render but also gives the best results.

Once the animation sequence is created as a wireframe model the next step is to render the animation as an .FLI or .FLC file. These are animation file formats. The .FLI is the original animation file format from Autodesk, and is limited to 320 x 200 size animations. The .FLC format is much more efficient, and also is flexible in image size. At all stages throughout this thesis .FLC files were used to store animations. Before rendering, the system first needs to be configured. The Renderer/Setup/Configure option displays a dialog box that allows a designer to set certain characteristics about the way that a animation is rendered. The option chosen for all of the animations in D-VR3 was "no compression", therefore, ensuring the highest quality output, and animations to be output as an .FLC file. There are three palette choices as indicated in Table 5.1. The type of palette chosen for use throughout this thesis was the High setting. This option was chosen because it produces the best results but, unfortunately, also produces extremely large files.

Finally, the last characteristic that needs to be set is the image size. Three parameters need to be set: width, height and aspect ratio. Width and height are obvious. The primary consideration when deciding on image size is the machine that the animation will be run upon. The smaller the image, the better it runs on less powerful machines. All the animations in D-VR3 were sized at 640 x 480. The aspect ratio is the ratio of pixel width to height. This ratio in the animations in D-VR3 was always set at 1.

After the configuration is set, it is time to render. To render chose option Render/Render from the option menu and then click the viewport that is to be rendered. The animations in D-VR3 were all rendered from a camera viewport, thus providing a full 3D animation from a camera perspective. These files were saved as .FLC files and can be played back using Autodesk Animation Player for Windows Version 1.10.

5.9. Conclusion to the Prototype Model.

The design of the prototype models was hugely important for many reasons. Firstly, it helped the designer to gain experience in using each of the packages used to construct the VEs. Secondly, the evaluation of the prototype model was a vital element of the design process, both internally and externally. Internal evaluation includes the evaluation of each of the packages used to create the prototype models in an attempt to ascertain which would be the most useful from a tourist perspective. The external evaluation was a necessary step in order to determine what exactly was required by the end-user at all levels. The prototype also was invaluable in designing an appropriate interface for the application and most importantly helped to gain interactive feedback from the end-user on the situation of the application. After each package was analysed and the most appropriate packages chosen, an actual model of a tourist location, the Monastic Enclosure at Glendalough, was constructed using each package.

5.10. The Actual Model - The Monastic Enclosure, Glendalough.

The construction of the prototype model was an invaluable experience as it helped to provide a firm understanding of the concepts and principles of developing VEs, it helped to visualise the problems associated with each of the engines analysed, and helped to overcome these problems. Once the prototype model was constructed to a satisfactory level the natural progression was to develop a model more suitable to the tourism industry. The ancient and ruined Monastic Enclosure at Glendalough, County Wicklow, was chosen as this model.

When attempting to simulate a real-life situation the first requirement is a complete and comprehensive set of architectural plans. The main priority in the case of the prototype model of the Computer Section was to get acquainted with the features of the software packages involved and, therefore, architectural correctness was not deemed a requirement at that stage. But when the Monastic Enclosure at Glendalough was decided upon as the actual model the first step was to obtain a complete set of architectural plans. Several sets of plans for the Monastic Enclosure were obtained, thanks to the cooperation of Mr. Willy Cummins, architect for the OPW. Unfortunately, none of the sets of plans obtained were complete and comprehensive. In many of the cases only some of the measurements were present but a complete set of measurements could be calculated from the combination of maps provided.

The next step that was taken prior to the construction of the Monastic Enclosure was a site visit. The main reasons for this visit were, firstly, to gain an over all impression of the site, and, secondly, to take photographs and video footage for inclusion in D-VR3. This site visit was made on Tuesday the 23rd of January 1996. This date was chosen because it was thought that there would be few visitors to the site on a weekday in January. This was indeed the case and, therefore, photographs and video footage could be taken with as little obstruction and distortion as possible. These photographs were then developed and scanned onto the computer using the same technique discussed in section 5.7.3. The video footage was also imported to the computer using the same technique addressed in section 5.7.3. As soon as the architectural drawings were complete and the photographs and video footage imported successfully the construction of models using the appropriate package could commence.

As was the case with the prototype model it was decided that the best way to describe the steps taken in developing the VEs is to describe them in accordance with the software types used. Again the first type of package that will be addressed are the visualisation and modelling packages.

5.11. Three-Dimensional Modelling and Visualisation Packages.

The construction of the prototype model was relatively straightforward with this type of package but in the case of the actual model, scaling had to be introduced. In order to describe the steps taken to produce the actual model one must begin by describing the steps taken to construct the basic geometry of the Virtual World.

5.11.1. Visualisation Geometry.

Due to the knowledge acquired from the construction of the prototype model, the construction of the Monastic Enclosure at Glendalough was relatively straightforward. The only difference from the prototype model was that this time the model was constructed to scale. This did not pose any great difficulty as Virtus Walkthrough supplies the designer with a wide range of measurements systems, Figure 5.25, to choose from. In the case of the buildings in Glendalough the buildings were modelled using imperial units mainly because the architectural drawing were measured using the imperial unit system.

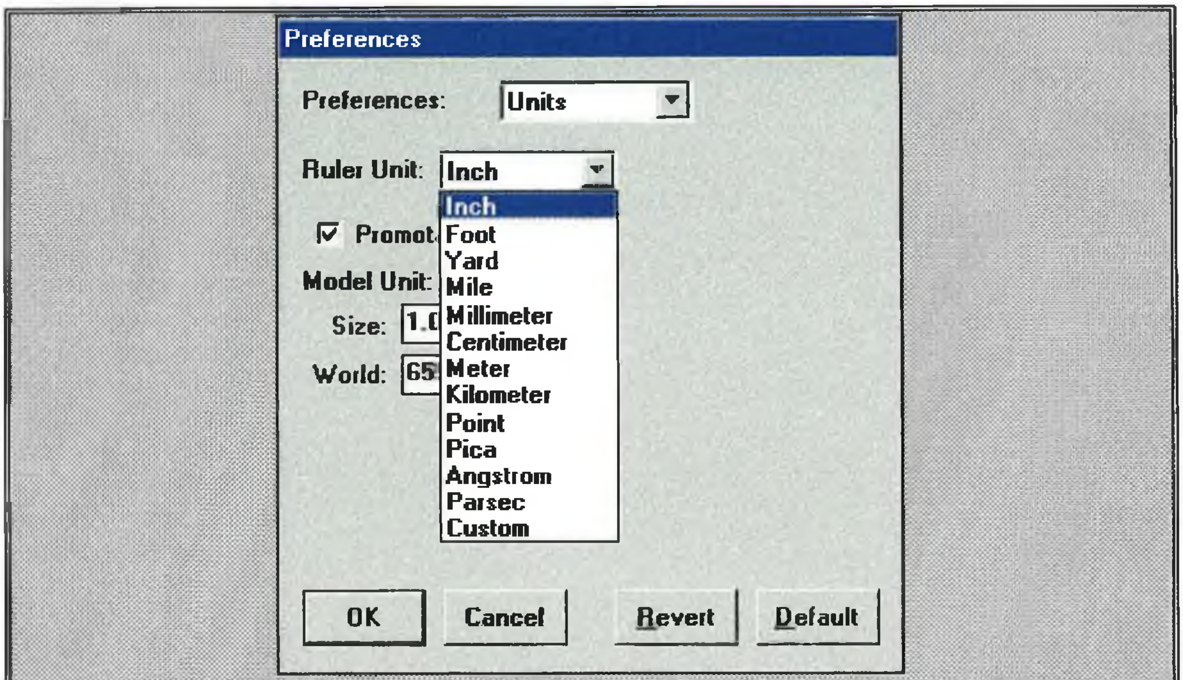

Figure 5.25. The Different Measuring Systems Provided in Virtus Walkthrough Pro 2.6. Beta.

Once the units to be used had been decided upon the next step was to construct the buildings. Each building was constructed, firstly, paying attention to the horizontal measurements of the buildings and, then, to the vertical measurements.

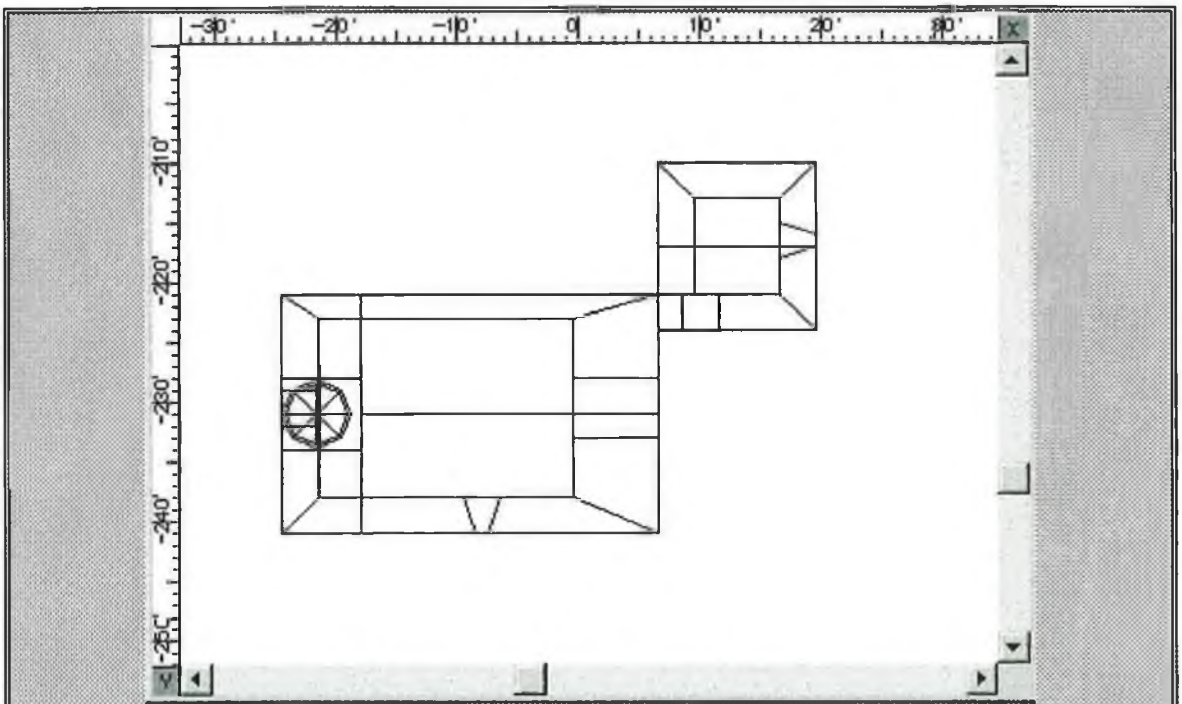
Figure 5.26. St. Kevin's Church During Construction Using Virtus Walkthrough's Top View.

The horizontal measurements were constructed exactly using Virtus Walkthrough's Top View, illustrated in Figure 5.26. The scale can be seen along the left side and the top of the view. The scale shown in this representation is the imperial scale, each unit displayed is one foot. In order to view this model in inches a designer would simply zoom in on the model. This scale is accurate to $^1/_{60}$ Inch.

Once the horizontal construction of the buildings was complete the designer's attentions then turned to the building's vertical construction. Each building's vertical construction was a lot more difficult that the horizontal construction. This was due to two main problems which arose in the case of some of the buildings. The first problem which arose was due to the fact that Virtus Walkthrough was not very well equipped to deal with arches. Instead of the designer being able to create a curve one actually had to create a series of lines. The main problem with this method of construction was that the more lines and polygons that the engine had to render the longer it took to redraw the scene as the participant moved through it. Therefore, the participant moved through the world at a much slower rate than when there was no arch in the scene. With the inclusion of a large number of arches, for the roofs and the windows, the participant's movement through the world slowed to an unacceptable level and, therefore, it was decided that a compromise had to be made.

The second problem that arose when dealing with the vertical construction of the buildings was that the surrounding topography had to be taken into consideration. It is almost impossible in a model such as this for a designer to take every mound and hollow into consideration but in early models of Glendalough, the larger topographical characteristics were built into the model. These were later removed because, yet again, the engine could not render these models fast enough due to the large amount of topographical detail present.

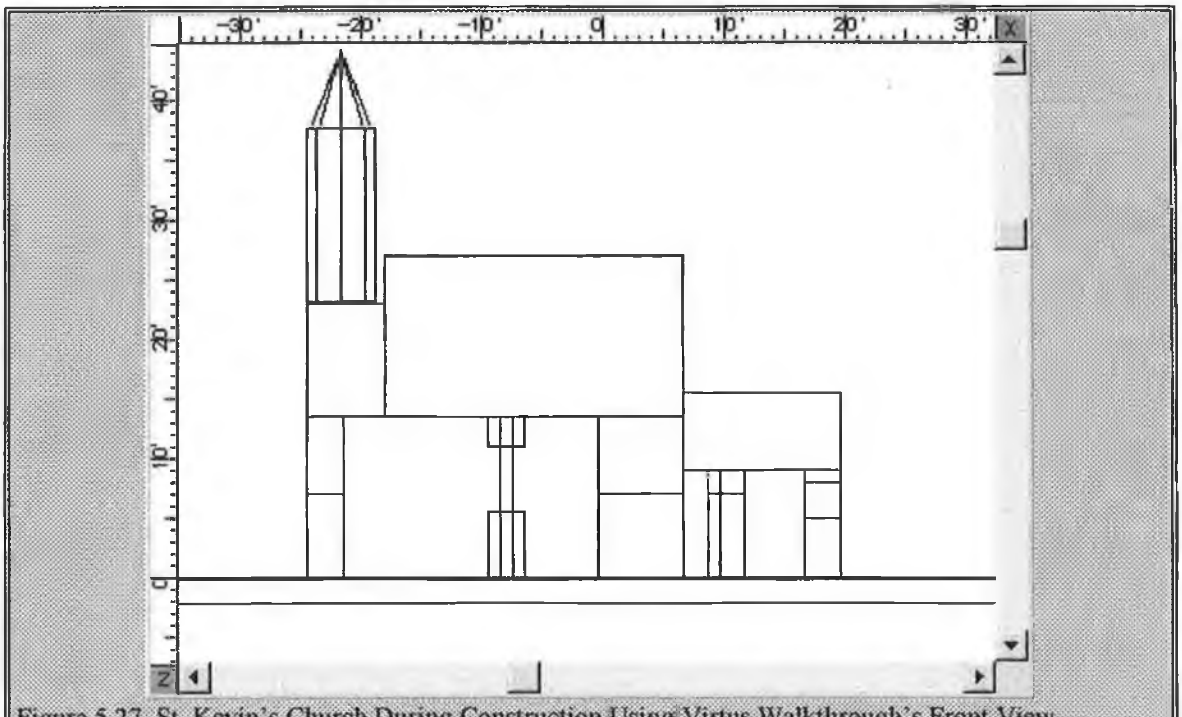

Figure 5.27. St. Kevin's Church During Construction Using Virtus Walkthrough's Front View.

Aside from the two problems which arose, the vertical construction was conducted, using any one of the Front, Back, Left or Right Views, in the same way as was the horizontal construction. All four of these views can be used at the one time along with the Top and Bottom View if required, but it was found that three views would suffice for even the most complex of model. The vertical construction is clearly illustrated in Figure 5.27.

Once the horizontal construction and the vertical construction were complete the model then stands as a wireframe representation, similar to the representation of St. Kevin's Church illustrated in Figure 5.28. This in essence is the skeleton of the building but in order to improve the building the designer would first of all have to apply shading and lighting, and then would have to apply an appropriate texture map. At this stage of the process the wireframe models of each of the seven buildings in the Monastic Enclosure was exported to 3D DXF files. These 3D DXF files could then be converted to most other modelling, animation or CAD packages. In the case of D-VR3 these .DXF files were in turn converted to the 3D Studio file format, .3DS, and the VRML file format, .WRL, for use on the Internet. These conversions are discussed more thoroughly in sections 5.13.1. and 5.14.2. respectively.

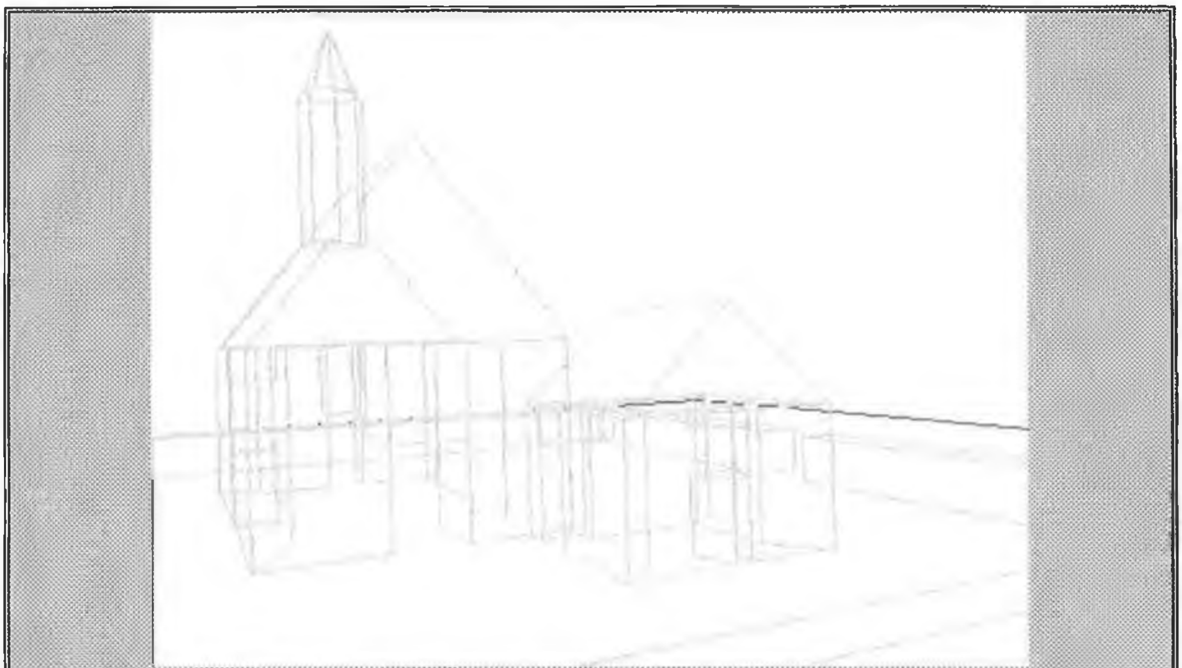

Figure 5.28. A Wireframe Representation of St. Kevin's Church, Glendalough.

5.11.2. Visualisation Shading and Lighting.

In some applications a world represented by a wireframe model might be perfectly acceptable, but in D-VR3 the main aim of the models created is to present a potential tourist with an impression of a tourist location. This would be very difficult to do with just a wireframe model; therefore, in this section the application of different types of shading and lighting attributes to the model is discussed. The first step that was taken in this process was the application of a simple white fill to each of the wireframe representations that were constructed in section 5.11.1. Although, this model still lacks realism, the application of a simple white fill, illustrated in Figure 5.29, supplied the designer and the participant with a far better impression of the buildings. The white fill is achieved by altering the shading and drawing options in the rendering properties in Virtus Walkthrough.

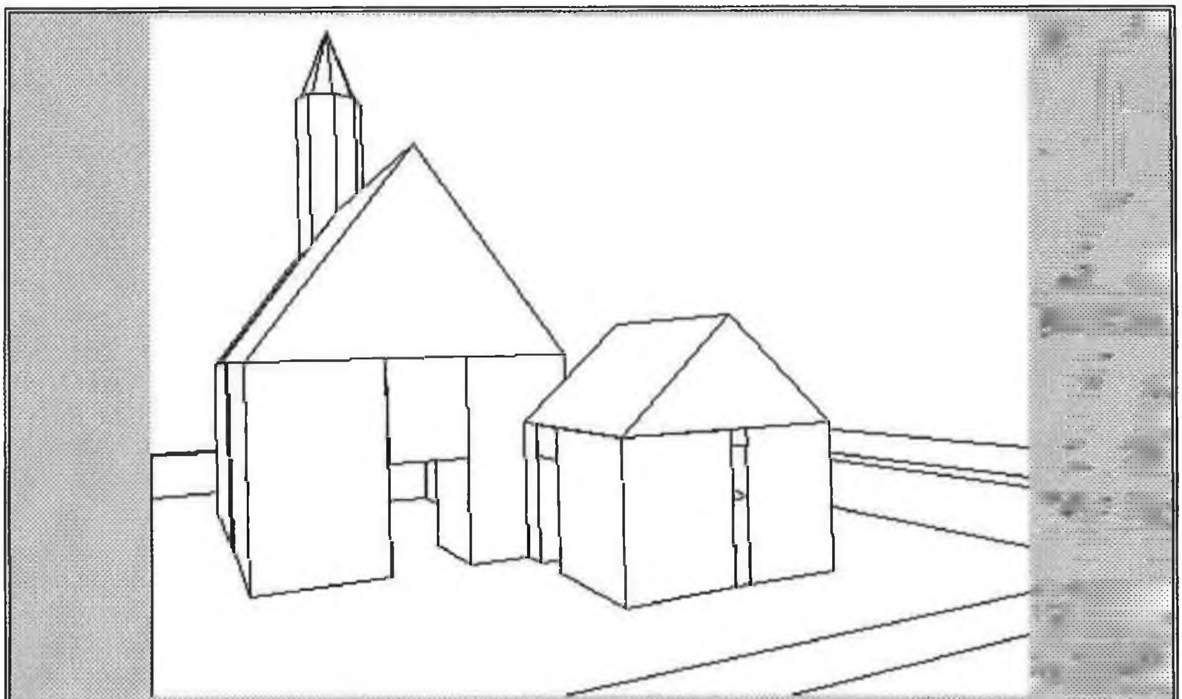

Figure 5.29. A White Fill Representation of St. Kevin's Church, Glendalough.

The next adjustments that were made to the models was the addition of an unshaded fill. A model of St. Kevin's Church in Glendalough which supports an unshaded fill is displayed in Figure 5.30. An unshaded fill in Virtus Walkthrough enables a designer to add colour to models for the first time but still does not allow him to shade the actual model. This type of fill is also achieved by changing the shading and drawing options in the rendering properties of the Virtual World.

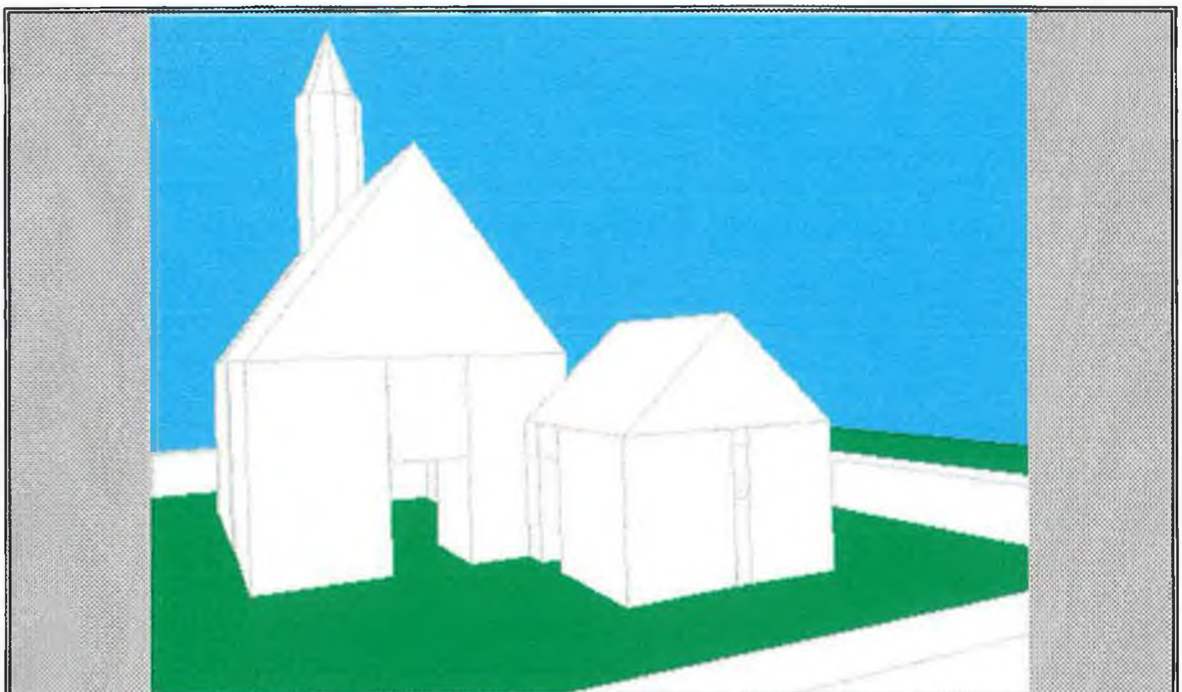

Figure 5.30. An Unshaded Representation of St. Kevin's Church, Glendalough.

The next step was the addition of a shaded fill. This enabled the designer to not only add colour to the each object but also to shade each colour. With this effect a designer can apply the first atmospheric attributes to the VEs. An example of a model which has a shaded fill applied is illustrated in Figure 5.31.

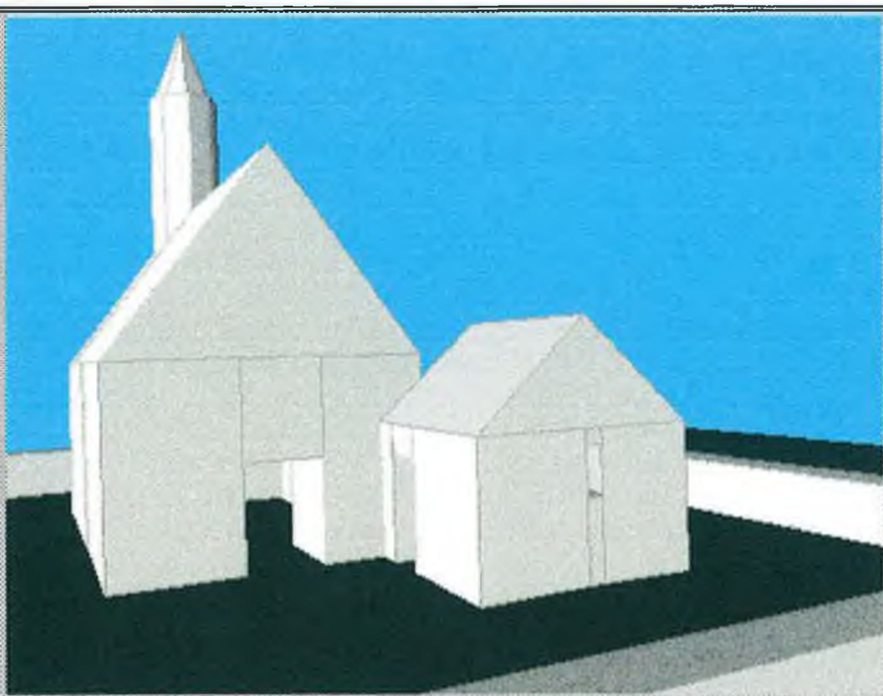

Figure 5.31. A Shaded Representation of St. Kevin's Church, Glendalough.

Lighting is a very important component of any VE, therefore, at this stage in the process lighting attributes were applied to each of the models. Virtus Walkthrough has its very own lighting editor which allows a designer to add directional or ambient light sources to an object or group of objects in a Virtual World. Differing degrees of lighting and shading were applied to each of the models in an attempt to provide as realistic an atmosphere as possible. Examples of the lighting effects applied to the model of St. Kevin's Church are displayed in Figure 5.32.

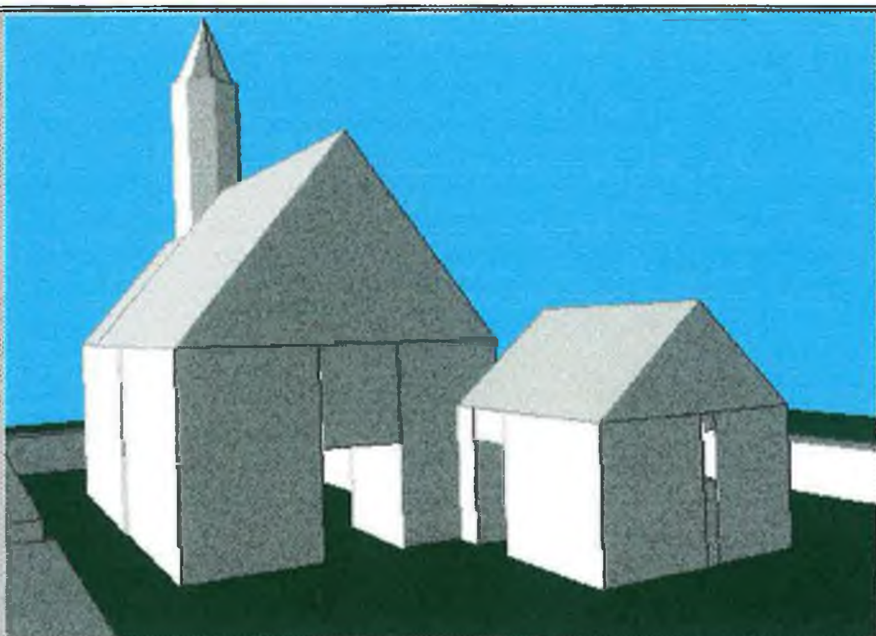

Figure 5.32. A Shaded Representation of St. Kevin's Church, Glendalough with Light Applied.

This stage of the design process has seen each building in the Monastic Enclosure in Glendalough progress from a wireframe representation to a well lit, fully shaded model. Although these models resembled the actual buildings in structure their appearance lacked realism. This was because of their lack of texture. For this reason the final step taken in the creation of appearance of each of the models was the application of an appropriate texture map to each of the buildings surfaces.

5.11.3. Visualisation Textures.

Computer graphics usually rely on texture mapping to obtain a realistic amount of visual complexity. Geometric models would be too large and inefficient for describing a complex scene in detail. Geometry is usually used only for the basic shapes and images are mapped onto the surfaces to add the necessary small variations in colour and orientation (Quarendon 1995).

In the case of the prototype model, texture was not applied to the model of the computer section because the packages used, Virtus Walkthrough Version 1.0. and Virtus Walkthrough Pro 2.0., did not provide an adequate texture-mapping facility, section 5.6.1. It was not until Virtus Walkthrough Pro. 2.6. (Beta Version) was acquired that an adequate texture-mapping facility was provided for the models created. This package provided a texture window, Figure 5.33, which enabled a designer to import textures of any size in both windows .BMP and .JPG formats. These textures, once imported, can then have their attributes changed to suit the task at hand. Textures can be applied to one face or a number of faces on an object by selecting the object and then select the texture to be applied from the list of textures already imported to the model.

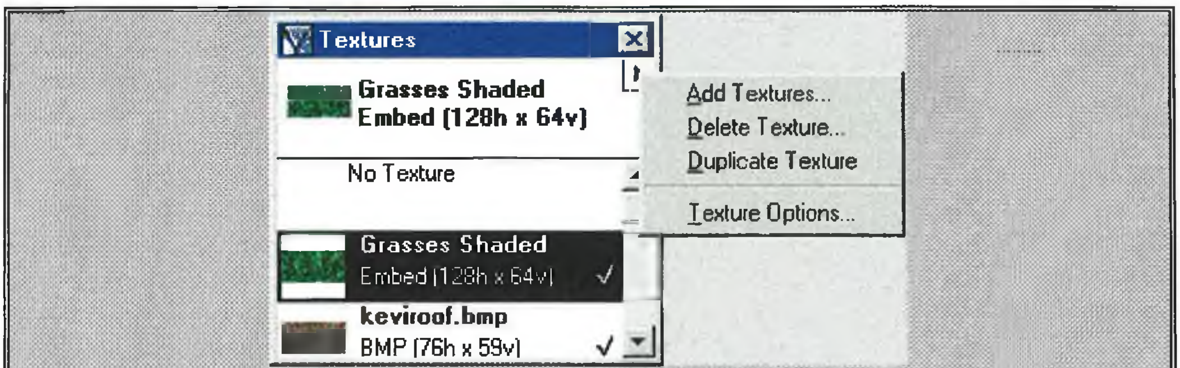

Figure 5.33. The Texture Window, Virtus Walkthrough Pro 2.6. Beta.

The textures applied to each of the models are the real textures taken from the actual buildings in Glendalough using the technique described in section 5.7.3. In certain cases the actual textures could not be obtained as the roofs or walls are no longer standing. In these cases the wall textures that were used were those taken from the most similar part of the same building and the roof textures used were taken from the roof of St. Kevin's Church which is still intact.

Figure 5.34. A Textured Representation of St. Kevin's Church, Glendalough.

Figure 5.34. illustrates the model of St. Kevin's Church with the appropriate textures applied. From this example it is difficult to see the quality of the texture maps that are applied; therefore, an additional example, Figure 5.35., provides a close-up display of these texture maps.

The addition of texture to each of the models of the buildings in Glendalough concluded the creation of appearance within the VE and the second stage which was dealt with was the addition of the Behavioural Characteristics to the environment.

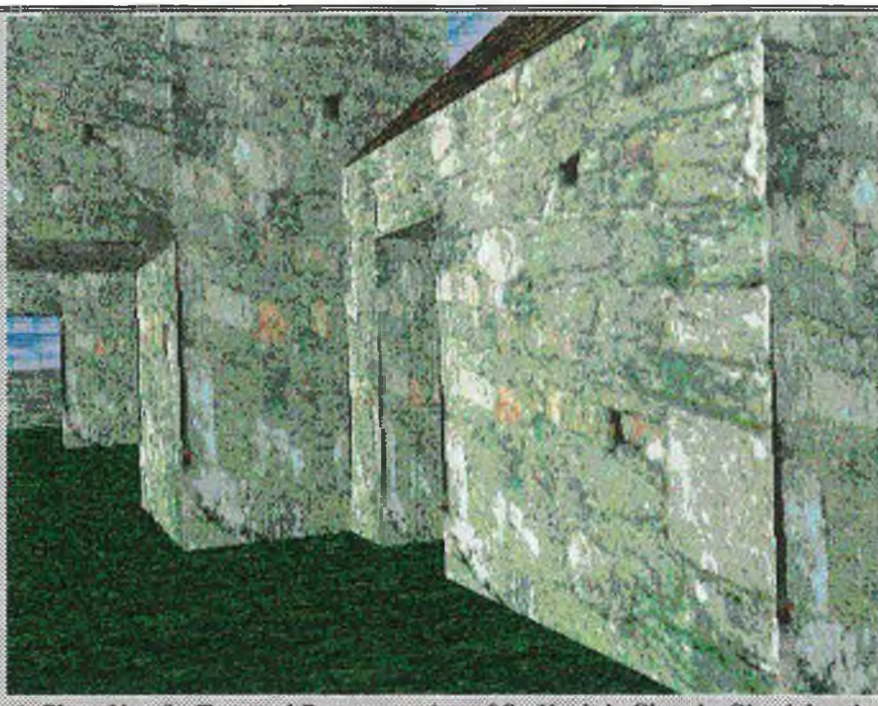

Figure 5.35. A Close-Up of a Textured Representation of St. Kevin's Church, Glendalough.

5.11.4. Visualisation Behavioural Characteristics.

As discussed in section 5.6.4., Virtus does not provide a designer with the opportunity to add many behavioural attributes to a VE. In fact, the only behavioural characteristic which Virtus Walkthrough does allow the designer to include is collision detection and this was only possible on the two later versions of Virtus Walkthrough that were used. The inclusion of collision detection is a much welcomed addition to both of these packages as navigation through worlds created with these products is a little difficult at first and participants have a tendency to pass through floors and walls, lose their sense of direction and become disoriented. The addition of the collision detection feature reduces the risk of this occurring and makes the navigation through such a world easier and more realistic.

5.12. Engines Used in Three-Dimensional Games.

As a direct result of constructing the prototype model using all three games engines it was decided to construct the actual model using the 3D Realms engine. The 3D Realms engine was chosen simply because it dealt with the execution tasks more easily, more quickly and, most importantly, better than the other two engines. In addition to this the 3D Realms engine supported some features that the other two engines did not. These features included stereo imaging, high resolution, and far more behavioural characteristics.

5.12.1. Games Geometry.

The main difference between the construction of the prototype model and the construction of the actual model was that the latter was constructed to scale. This caused some problems due to the fact that the 3D Realms engine did not include a recognised measurement system. The measurements system employed by this engine is a little confusing and slightly unorthodox. Everything is measured in units and the units are different horizontally and vertically. This creates the problem of equating the games engine's measurement units to real world units. The following scale was reached through equating wall heights and widths to known heights and widths such as the participant's dimensions. The participant's dimensions are 56 units high and 32 units wide. After many experiments the following scale was reached. In these engines, 16 horizontal units are equal to one horizontal foot in real-world dimensions and 10 vertical units are equal to one vertical foot in real-world dimensions. From these dimensions the participant's dimensions are 5 feet 7.2 inches high and 2 feet wide in real-world dimensions. When these conclusions were reached a spreadsheet model was constructed that would automatically convert real world dimensions to games engine's units. It is essential for a designer trying to simulate a real-world situation to have sound foundation and an accurate scaling and measurement system. Once a satisfactory measurement-conversion system had been drawn up the next step in the design process was to create the basic structure of the world by creating the world geometry.

As mentioned previously, the techniques used to construct VEs using the 3D Realms engine are slightly different than the techniques used in the Raven and Id engines. The first difference is to do with the tools used. Instead of using a 2D map editor and then having to view the world from inside the games engine, Build, a 3D Realms map-editor, enables a designer to view the world in the actual editor in either 2D or 3D mode. This saves a lot of time and effort when constructing the actual VE. The 2D mode, displayed on the left side of Figure 5.36., is quite similar to those used with the Id and Raven engines, but it is in the 3D mode, on the right side of Figure 5.36, where Build really excels. With the editor, any changes made in the 2D mode can be observed instantly in the 3D mode by pressing the [ENTER] key. The 2D mode is where the basic structure of the models is constructed and scaled and the 3D mode is where the heights of all the buildings are set and all the texture selection and manipulation is conducted. Certain constructions such as slope creation and sprite placement can also be carried out in the 3D mode. The Build editor has part of the 3D Realms engine programmed into it and this is how the editor enables a designer to view the environment in 3D mode.

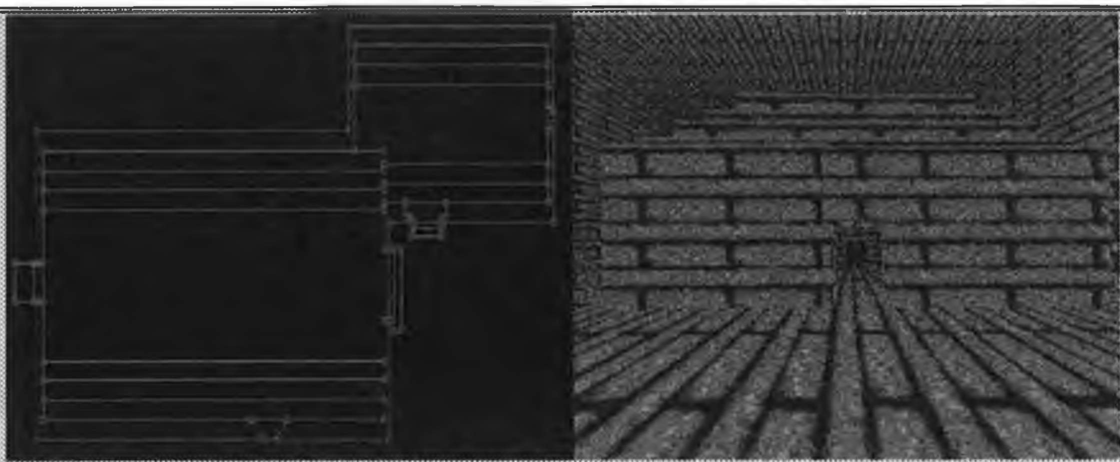

Figure 5.36. The 2D Design View, Left, and the 3D View, Right, from the 3D Realms Engine.

The first step of the design process was to construct the basic structure of each of the models. This was conducted mainly in 2D mode and then the height of each building was set in 3D mode. The first draft of each model was not constructed to scale and each building displayed the 3D Realms engine's default texture, illustrated in Figure 5.37. The reason that the building was not originally modelled to scale was that the building should first of all be constructed to determine whether the general layout was possible and once this layout was deemed satisfactory then the building would be scaled.

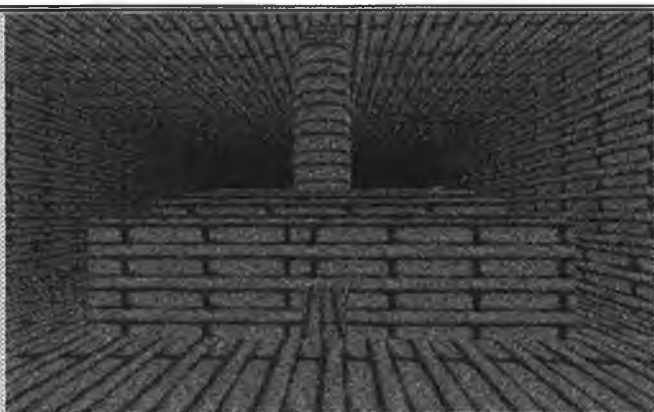

Figure 5.37. St. Kevin's Church Constructed Using the 3D Realms Engine (Default Textures & No Scale).

The difference in scale is very obvious if one was to compare the non-scaled model, Figure 5.37., to the scaled model, Figure 5.38. The scaled model of St. Kevin's Church also has a sloped roof. This capability was a welcomed addition as the 3D Realms engine was the first games engine that was examined that was capable of creating slopes. With all the other games engines roofs had to be constructed using a step approach. The roof of the Round Tower was the only building constructed using the 3D Realms engine whose roof was constructed using a step approach.

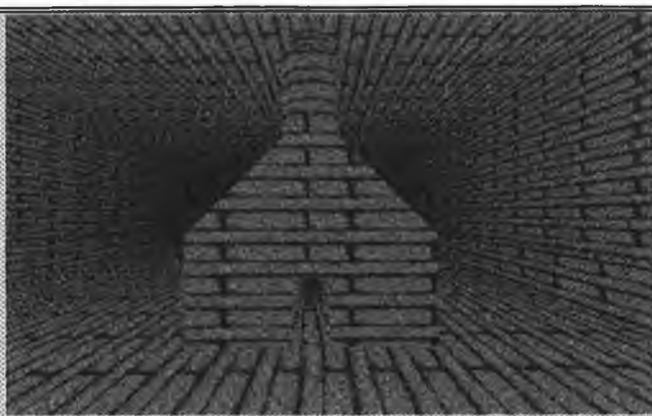

Figure 5.38. St. Kevin's Church Constructed Using the 3D Realms Engine (Default Textures With Scale).

Despite the fact that the 3D Realms engine was a far superior VR engine to the other two games engines examined, the 3D Realms engine still had certain limitations. Although the 3D Realms engine does allow one level to be built above another, it does so with one caveat which severely restricts the usefulness of the facility. This is that when one level is placed above another the two levels can not be seen at the one time. This vastly reduces the number of situations in which the feature can be used. It can not, for example, be used for the construction of roofs which would be the principle place that such a feature would have been used. In order to create a roof on a level which could be observed by a participant two other techniques were attempted. The first technique involved the placement of several sprites to simulate a roof on top of an open sector. This technique would work perfectly on a building which required a flat roof but all the buildings in the Monastic Enclosure in Glendalough all had sloped roofs of one type or another. This technique unfortunately would not be suitable for the creation of roofs. The second technique examined is called the portal approach. Basically, the concept is of two worlds being designed identically, or one world being designed and then copied to another location. The first world allows the participant to view the outside of the building with a roof and all the outside textures and the second world allows the participant to view the inside of the building with all the attributes that it may possess but the only way a participant may view both the outside and the inside of the building is with the use of a portal. A portal is a device, usually located at a doorway or window, which allows the participant to be transported from one world to the other. The concept behind a portal is that once a participant enters a portal, he immediately materialises at some other pre-determined location. There are three components required to construct a portal. They are:

A portal trigger line.	This is the line that a participant must step over to enter the portal and this triggers the portal action. This is achieved by applying a special line type to this line. Unlike other special line type this trigger can only be activated from one side. This is logical as otherwise as soon as a participant tries to leave a portal they would be transported back to the original portal.

A destination sector.	This is where the participant is materialised.
A landing Spot.	This marks the precise location within the destination sector where the participant is materialised. The 3D Realms engines requires more information than just the sector in which the participant will be transported to, it needs to know the precise location. This is marked by placing a marker within the destination sector to signify the portal's destination. Probably the most interesting attribute of a landing spot is its facing-angle attribute. This attribute was put to use in the 3D Realms version of Glendalough to face the participant a specific direction when they emerge from a portal.

There are certain attributes which are usually used in conjunction with a portal and, therefore, make it identifiable. These attributes include such items as specific textures on the ground of the portal, a step into the portal sector and a specific sound and flash which signifies the participant being transported to a different location. In the case of D-VR3, these attributes have all been removed in order to move the participant from one location to another without him being aware of such a movement, thus, allowing the participant to believe that he is moving into a model and not being transported to a different model. This approach has been used to good effect in the construction of the buildings in Glendalough using the 3D Realms engine.

5.12.2. Games Shading and Lighting.

Once the basic structure of each of the models had been constructed using the 3D Realms engine the next step was to add appropriate lighting and shading to them. As discussed previously, section 5.7.2., the 3D Realms engine has its own default shading which it automatically adds to each model. An example of the model of St. Kevin's Church in Glendalough is clearly illustrated in Figure 5.28. Therefore, the only process which needs further discussion at this point is the addition of lighting to each of the models. The addition of lighting to models in the 3D Realms, previously discussed in section 5.7.2., is a simple, if somewhat laborious procedure which involves the application of lighting to each of the surfaces within a sector. Although this may be a time consuming process it is one that allows the designer total control over the lighting within a world.

5.12.3. Games Textures.

This section describes the addition of actual textures taken from the buildings in Glendalough to each of the models designed and the removal, or in many cases, replacement of textures in an attempt to make the interface more acceptable to the user. The first process dealt with was the inclusion of the appropriate textures to each of the models. The first step of this process was the conversion of the photographs taken into a suitable format for placement into the 3D Realms engine. This procedure is outlined in section 5.8.3.3. Once the textures were in the correct format, their actual inclusion into the 3D Realms engine using Editart was a relatively straightforward process. There are, however, certain factors which must be given some consideration.

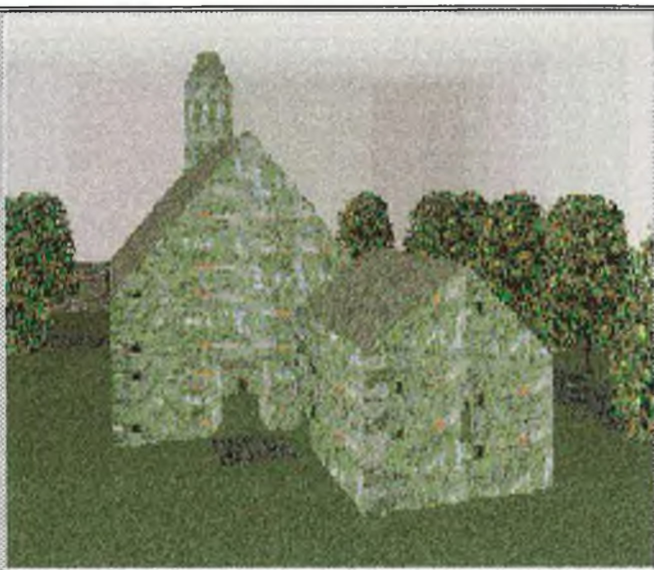

Figure 5.39. St. Kevin's Church Constructed Using the 3D Realms Engine (Actual Textures With Scale).

The 3D Realms engine allows for a total of 4095 graphic images to be used. Of these 3585 are actual graphics that are used by the game Duke Nukem 3D and the remaining 510 spaces are specifically provided for designers to import graphics files. Designers can also import graphics in place of any of the original 3585 graphics. Therefore, a designer can use a total of 4095 textures in any one VE. In some cases these textures are not all single textures: they may be animated textures or textures linked together to create characters. In the case of one animated character a total of 109 textures would have to be imported if one wanted to replace the character. In the case of both the prototype model and the actual models, the total number of textures used was 456. All the textures used in the prototype model and in the actual models were created or scanned from the actual buildings using the technique described in section 5.7.3.

Once the textures had been imported into the 3D Realms engine, it was then possible to apply them to the appropriate surfaces using Build. In most cases the texture had to be shaded, resized, aligned and/or oriented relative to an adjacent wall. These functions were conducted in the editor's 3D mode. In some cases an object's surface would have to be divided into smaller sections to accommodate the application of some textures. A large amount of time was spent manipulating the textures but it is essential in order to create a realistic VE. Figure 5.39. illustrates a model of St. Kevin's Church, Glendalough, complete with its actual textures.

In addition to the creation of Virtual Worlds using the 3D Realms engines, the actual appearance of the interface itself also caused problems. If these environments were to be used for tourist purposes certain parts of the interface would have to be changed. For instance, the introduction to the game provides the player with a list of startup commands, a loading sequence, a Main Menu and also provides sound effects to accompany each of the above and this is even before a player gets to play the game. Once the player enters the game the engine then provides him with an information bar at the bottom of the screen which provides information on the player's health, the types of weapons the player has, and the amount of ammunition the player has for each weapon. The screen also display the player's hand with the current weapon in their hand. All of these items mentioned have to be removed or at least replaced in order to make such an engine suitable for use as a tourist information tool. All these screens and sounds were changed accordingly in order to make the interface more suitable for the tourism industry; see Appendix H for more detail.

5.12.4. Games Behavioural Characteristics.

The 3D Realms engine offers a designer a choice of many different behavioural attributes but in reality there is very little need for much behaviour in the Monastic Enclosure in Glendalough. These attributes enable the designer to set several different behavioural situations such as elevators, portals and a host of lighting effects. Portals were the only behavioural characteristics used in the worlds and they were used to speed up the models. The portal approach, previously described in section 4.5.2. and 5.12.1., was a necessary step to take because otherwise the computer would have to redraw the whole Monastic Enclosure, rather than just the part that the participant is in at that point in time, every time the participant moved and this would slow down the participant's interaction rate to an unacceptable speed. A second reason for using the portal approach was that otherwise the "sector above a sector" effect could not have been achieved. Portals were placed at the doorways and windows to each building to transport the participant to the inside or outside of the building depending on his current position. Portals were also used to transport a participant from one world to another. Portals were placed at the entrances and exits to each world so that if a participant steps into a portal he is immediately transported to the entrance of another world. This process was conducted in such a way as to create the appearance of a participant stepping through a doorway into another room.

The 3D Realms engine also makes use of computer operated characters. These help to populate a world in an attempt to increase one's sense of realism and scale within an environment. Wall textures are relatively straight-forward to create and import once the specifications are met but placing another participant or computer operated character into the engine is an extremely time consuming exercise. This is because wall textures are generally constructed from just one texture whereas in the case of a new participant or a computer operated character, a texture has to be created for each of the character's movements from eight different angles; front, back, left, right and the four diagonal angles. In the case of the Monastic Enclosure, all the computer-operated characters took the form of medieval monks which each required 52 separate images to be replaced. An example of these computer operated characters can be seen in Figure 5.40.

The 3D Realms engine, as well as all the other games engines examined, supports collision detection. This is not unusual because a game engine of this type which did not support collision detection would be relatively useless.

Figure 5.40. St. Kevin's Church Complete with Computer-Operated Characters.

5.13. Three-Dimensional Graphics and Animation Packages.

The construction of the prototype model using 3D Studio helped to master the features of the animation package and to become familiar with the concepts of constructing high quality animations. The method used in the construction of the actual model, which follows, was extremely similar to that used in the construction of the prototype model described in section 5.8.

5.13.1. Animation Geometry.

Three-dimensional graphics and animation packages excel in the creation of complex geometry and often are able to describe the interconnections of various complex geometric parts at the one time. The output of a typical animation package often consists of a wireframe description of the scene's geometry, Figure 5.41. In the case of D-VR3, the geometry was imported in 3D DXF format from the models created using Virtus Walkthrough, described in section 5.11.1. Although wireframe representations are not very appealing, this is the skeleton on which one can build a VE.

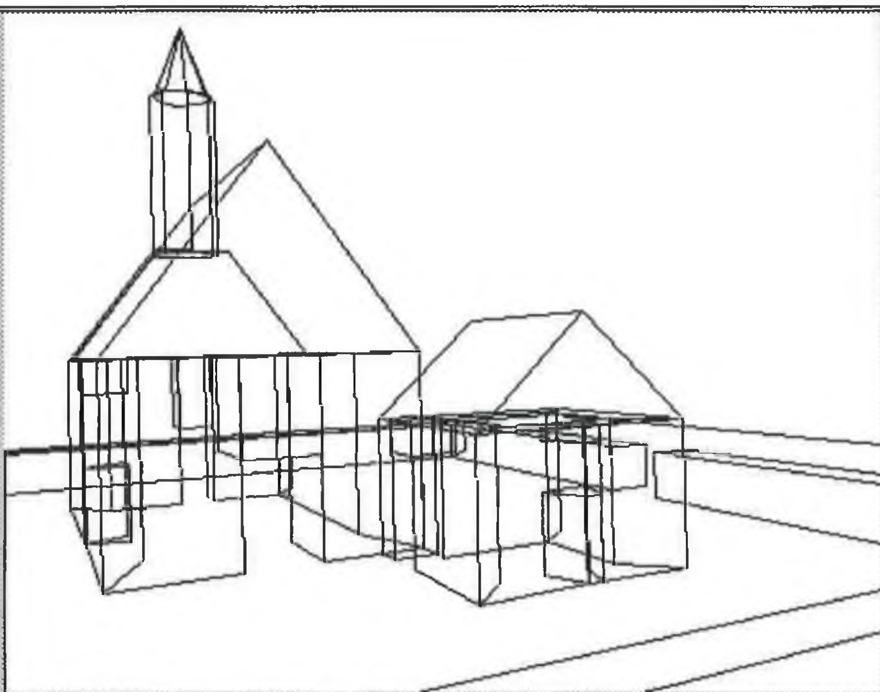

Figure 5.41. A Wireframe Representation of St. Kevin's Church Imported from Virtus Walkthrough.

After the basic model had been imported from Virtus Walkthrough, the next step involved the inclusion of camera positions in order for the model to be viewed from a specific position. Camera viewpoints are added using the Camera/Create option from the command menu. A number of camera viewpoints can be used to heighten interest and to enable the participant to view the world from several different angles but in the case of the models of the buildings in Glendalough only one viewpoint was used. Because the purpose of the exercise was to create an animation which would bring the viewer on a walkthrough of each of the models. For this to be achieved only one viewpoint was necessary.

5.13.2. Animation Shading and Lighting.

Once the camera viewpoint was set and all of the objects were created and located in appropriate positions in the world, the next step was to add shading and lighting in order to increase the realism of the world. Lighting affects mood and emphasis. Rarely, if ever, does a scene depend on just one light source, even when a scene is supposed to be simulating light from the sun or a single light bulb. A complex configuration may be needed to convey the most simple of effects. First, the lights are positioned to examine the effects that it has on the environment. Then, for each light, the type of light source is chosen. A light can be ambient, omni or spot. A designer can add a number of different effects to create a warmer or cooler effect. A shadow may even be added. The addition of shading and lighting was relatively straightforward. A Grey texture was chosen from the libraries provided and applied to each of the models using the Surface/Material/Choose and the Surface/Material/Assign/Object commands respectively. A simple ambient light source was added using the Lights/Omni/Create option from the menu. Once lighting and shading was added the model was then rendered, Figure 5.42. A spot light was then added to the models using the Lights/Spot/Create options from the command menu in an attempt to simulate the appearance of the sun.

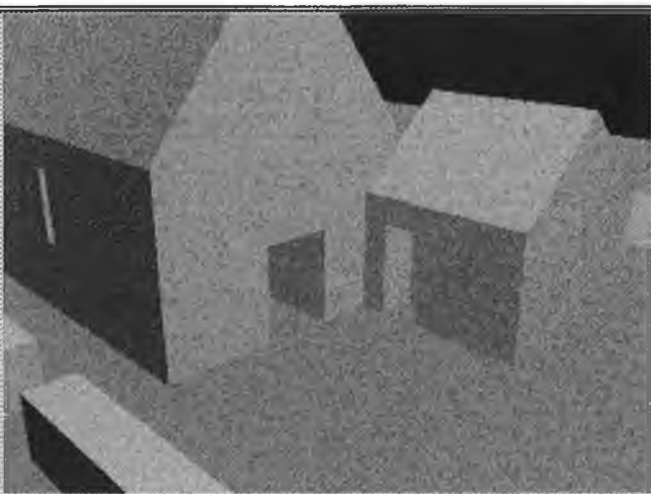

Figure 5.42. St. Kevin's Church, Glendalough, with Shading and Lighting Applied.

5.13.3. Animation Textures.

The application of texture is an extremely important part of any VE and often makes all the difference in relation to the realism of the world. Many 3D modelling and animation packages contain a list of materials that can be applied to the objects but in order to make the scene more realistic additional textures can be created or imported from real-life objects. The first step that was taken when applying the actual textures was to map the coordinates of each texture onto the surface of the appropriate objects. This was conducted using the Surface/Mapping/Aspect/Bitmap Fit option in the command menu. Figure 5.43. is an example of St. Kevin's Church with a texture applied which is not mapped correctly. The next step taken was to choose the actual textures using the Surface/Material/Choose command and then apply them to the appropriate models using the Surface/Material/Assign/Object option from the command menu. Figure 5.44. is an illustration of the final model with the textures mapped correctly.

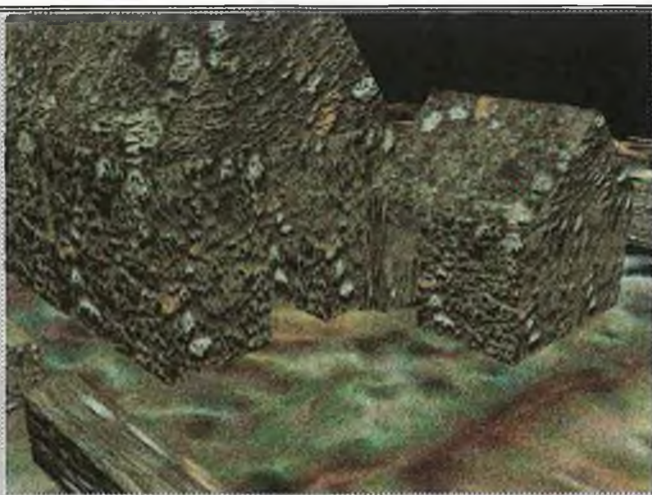

Figure 5.43. St. Kevin's Church, Glendalough, with Textures Applied But Not Mapped Correctly.

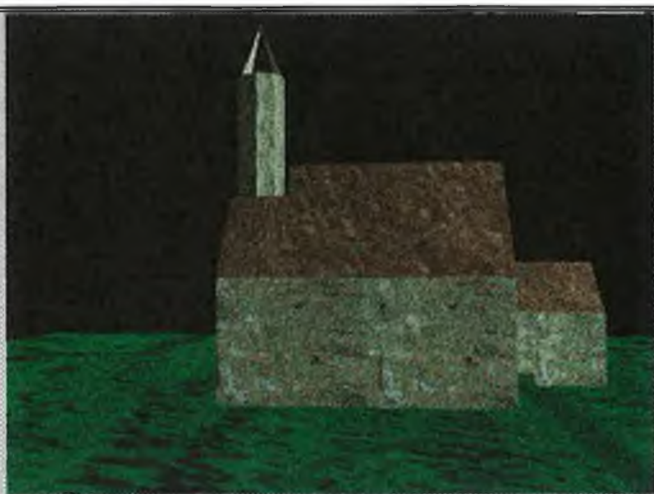

Figure 5.44. St. Kevin's Church, Glendalough, with All Textures Applied Mapped Correctly.

5.13.4. Animation Behavioural Characteristics.

The process thus far has been associated with the development of one still image. In certain circumstances this might be all that is required but in the case of the creation of a preprogrammed walkthrough of a VE an animation sequence is necessary. To create an animation sequence, certain *key frames* for each of the pivotal moments within the world had to be created. Once the key frames were created 3D Studio was used to fill in all the other frames between these key frames. This was achieved by first deciding on the number of frames that were going to be used in each animation, then selecting the objects that were to be moved and describing their movement and specifying how many frames the movement should take. An object can be moved along a line, a trajectory or in a circle; they can also be rotated, scaled, squashed, or twisted.

The objects in the scene are not the only elements that are subject to movement. The camera and lights may also be animated. The camera can be made to zoom in and out, pan from side to side, or be hoisted up from eye level to sky level for an aerial view. This was the technique that was used most often in the creation of the animated sequences of each of the model in the Monastic Enclosure in Glendalough. Likewise, the lights can be made to move. Using this technique the rising and setting of the sun can be simulated. Finally, all the information pertaining to the scene is rendered in order to produce an animation.

5.14. Other Packages Considered.

There were several other software packages considered for creating VEs. Two of the most useful were Apple's QuickTime VR and the VRML.

5.14.1. QuickTime VR.

QuickTime VR enables a participant to create a VE using stills taken with a panoramic camera or a 35mm camera. These stills would then be digitised using the technique outlined in section 5.7.3. The images would then have to be gelled together to form a 360 degree scene in a program known as the stitcher. QuickTime VR also enables a designer to include hot-spots into the world that perform a specific action once they are chosen. The action may be to display text, play a sound or most commonly to open another world. The participant is allowed total freedom to navigate through these worlds using an appropriate control device. A participant can zoom in on a specific object for a closer look, or zoom out to see the big picture. Apple calls a Virtual World that provides this freedom of navigation a panoramic movie. QuickTime VR movies can only be created using the QuickTime VR development kit on an Apple Macintosh computer. However, these movies may be viewed using a run-time version on other platforms. The development kit requires a high level of knowledge and a significant amount of programming. QuickTime VR has good potential for the creation of Virtual Worlds but certain factors such as the cost of the hardware and software required prohibited its use in D-VR3. This option was not pursued any further due to the sheer expense of the equipment required in its implementation.

5.14.2. Virtual Reality Modelling Language.

VRML is a language for describing multi-participant interactive platform independent VEs networked via the Internet and hyper-linked with the World Wide Web. All aspects of Virtual World display, interaction and inter-networking can be specified using VRML. VRML allows VEs to be incorporated into the World Wide Web, thereby allowing users to explore a Virtual World on the Internet.

There are currently two versions of VRML available Version 1.0. and Version 2.0., alternatively known as Moving Worlds. The first version of VRML allows for the creation of Virtual Worlds with limited interactive behaviour. These worlds can contain objects which have hyperlinks to other worlds or to other data types, in which case the client must launch the appropriate viewer. Although the introduction of VRML V1.0. was quite an accomplishment, it does not specify a way to include animation, allow for sound, direct interaction with objects, or real-time interaction with other people. VRML V2.0, was developed in an attempt to address these issues. None of the proposals have explicit support for multiple-user environments, though all make such support easier. The constraints faced by 3D world designers differ too widely to "hardcode" multi-user support into the standard. For this reason, the inclusion of multiple-user capabilities into a VRML world is possible but other hardware is usually required.

With the introduction of VRML, a number of companies have introduced VRML-related products. Most tools today fall into one of the following categories: 3D modelling tools with the ability to export VRML; VRML authoring packages; and VRML browsers, which are used to view these worlds. In the case of D-VR3, the VRML worlds were directly exported from worlds previously created in Virtus Walkthrough, as described in section 5.11. During this exporting of the worlds, the texture maps, which had been stored as Windows .BMP files, were converted to a .JPG file format in order to reduce the size of the VRML file and, therefore, the download time required to view the file over the Internet. The worlds created were all VRML Version 1.0. but can easily be converted to Version 2.0. As one would expect the VRML world, illustrated in Figure 5.45., is extremely similar to a world created using Virtus Walkthrough, Figure 5.34.

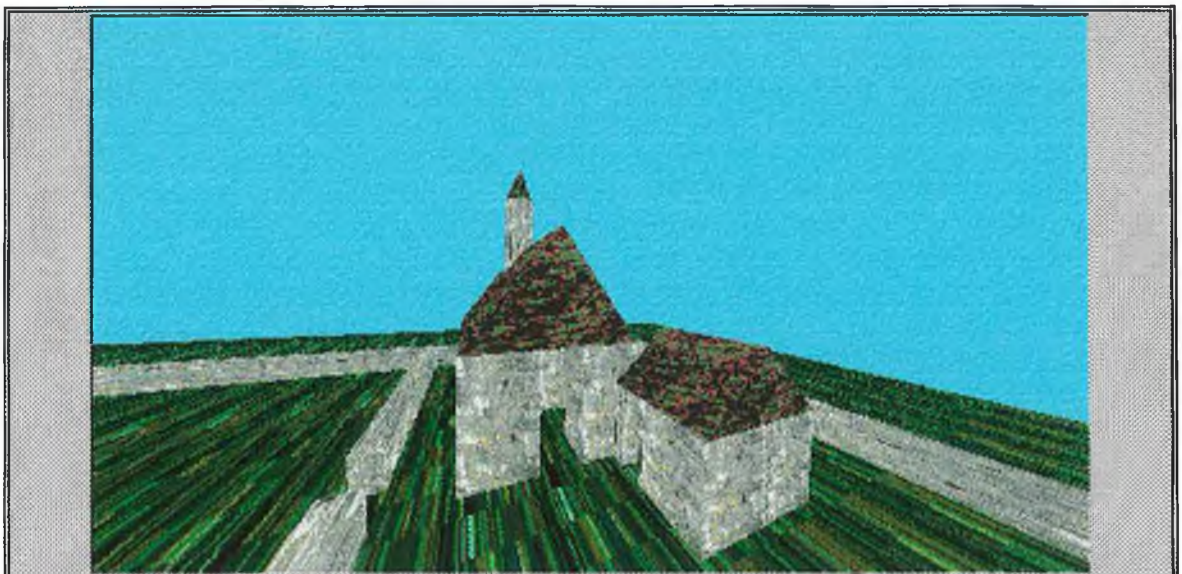

Figure 5.45. A Walkthrough of St. Kevin's Church Using the Virtual Reality Modelling Language

Once the worlds are created, they can then be viewed with an appropriate browser either over the Internet or on the host machine. A VRML browser interprets the scene descriptions and renders the resulting images on the host computer. Rendering is performed through the lens of a virtual camera that moves, tilts, and twirls in response to user input. Many firms are developing extensions to existing Internet browsers that provide VRML capabilities as enhancements or plug-ins. These products enable users to access VRML worlds with their current browser technologies. The browser that was used throughout the course of this research was a Netscape plug-in called Live3D. This browser was chosen because it allowed a participant to navigate through the VRML worlds in a fast intuitive way. The potential for such an application to the tourism industry is enormous. A potential tourist could examine tourist locations, Hotel rooms, conference facilities, or banqueting facilities in great detail over the Internet prior to a trip in order to aid them in their travel decisions.

5.15. Other Factors that Must be Taken into Consideration.

Apart from the steps that were taken in order to build each of the environments outlined above, there are other considerations that must be discussed when experiencing a VE. These considerations such as resolution, stereo imaging, navigation and conversion standards are addressed in detail in Appendix I.

The next stage in this thesis was to place the VEs created into a user-friendly intuitive interface in order to make them easily accessible to the user. In order to achieve this it was decided to use a multimedia-based environment. As this is only a secondary aim of the research, it is described only briefly.

5.16. Interface Design.

The interface is the means through which a participant can communicate with the environment, the interface devices, and the process by which a designer can store the environment, the multimedia application. The first step to be addressed is the participant's means of communication with the VE. Due to the focus of this research being placed on the use of VR on desktop computers the manipulation devices are simple: generic mice and keyboards. D-VR3 can also be configured to handle joysticks, 3D mice, DataGloves, and HMD's.

The second step in the design of the interface was the construction of a multimedia application in order to join the VEs created with the traditional types of tourist information. Creating a multimedia structure/module for any application can be divided into four steps:

1. Capture or create content;
2. Edit content;
3. Combine elements and add interactivity;
4. Deployment.

Each step is now discussed:

1. The first step is to create or capture digital content from an analog source. The digital content of a multimedia application can be comprised of text, images, animation, audio and video. In the case of D-VR3 all these types of contents were used plus the addition of VEs. Several different techniques were used to add these contents to the multimedia applications. In the case of the text, it was either scanned using an optical character reader or typed in from different historical and architectural books about Glendalough. The images were all scanned in using the technique described in section 5.7.3. The animations were created using the technique described in section 5.13. The audio files were created using Creative Wave Studio and the video was captured using Asymetrix DVP. Finally, the content also included the VEs described in sections 5.11. and 5.12.

2. Once all the content had been added the next step taken was to edit the content. This step mainly involved the resizing of the images and editing the imported videos. The images were resized in order to fit into a specific place within the Multimedia application and these modifications were carried out using PhotoPaint Version 6.0. The main reason for editing the video clips was because

they were just too large to handle effectively. The editing was conducted in DVP. The remaining contents did not require editing.

3. The third step was to combine all the elements and add interactivity. It was at this stage that the Multimedia tool had to chosen. This was a very important decision to have to make and, therefore, a number of different packages were considered. From all the packages considered, two packages were chosen. These two products varied in capability and ease of use. Both of the products examined neatly serve a different purpose and developer level. Gold Disk's Astound 2.0 is a sophisticated presentation package that emphasises multimedia capabilities, and creates standalone interactive applications whereas, on the other hand, Asymetrix's Multimedia Toolbook 3.0 offers an excellent combination of ease-of-use and power. Both of these packages were used to construct the application to house the prototype models but from these early experimentations Toolbook was chosen for the actual application. Asymetrix Toolbook was used because besides the fact that it had far more features than Astound it also handled jobs more cleanly, faster, more smoothly and produced more professional results. Toolbook also contains a scripting language called OpenScript that allows a designer complete control over the look and feel of their interface.

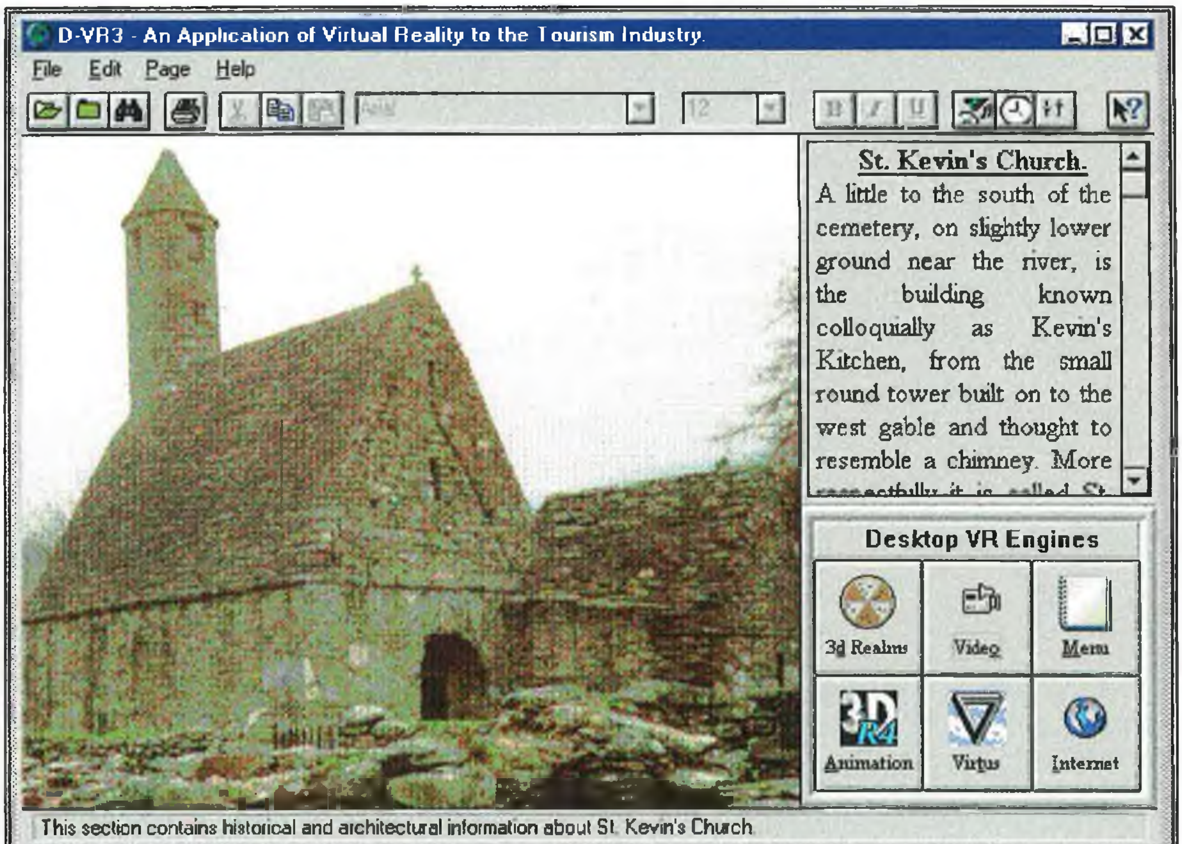

Figure 5.46. The Menu to St. Kevin's Church Glendalough.

D-VR3 is created as a structure in which different forms of tourist information can be inset into the one program. This structure is extremely flexible and allows for different languages and different media - past, present and future - to be integrated. The last factor that was considered was the issue of how each VE would be integrated into the multimedia-based environment. This is a very important issue as it makes the difference between whether an application is intuitive or not. In the case of all the Virtus Walkthrough files, .WLK, .WTP and .VVR, and the animation files, .FLC files, the extensions were added into the Windows initialisation file, WIN.INI, and the worlds were called up using Toolbook's scripting language. In the case of all of the worlds created on the games engines a Batch file was created in DOS to run the VEs with the necessary command line parameters. This Batch file was in turn executed by using Toolbook's scripting language. Using these procedures a participant only has to choose a button, illustrated in Figure 5.46., to load a specific world and no knowledge of the inner workings of the application is necessary.

4. Finally, the deployment of the application. In general, it's best to develop on the same platform that that the application will be delivered. If multi-platform deployment is a concern, developers will find some, but not many, options available. Not all authoring techniques transfer well from one platform to the next. A traditional marketing approach requires a 486/33 with VGA only. In that case, one would have to capture images at 640 X 480 with 256 colours. This ensures the application will run on any IBM compatible PC. If one can control the playback platform, one can increase the resolution and colour depth and greatly enhance the content. In an application, such as D-VR3, that contains video images and animations and virtual walkthroughs, increasing the colour depth to 64KB, for example, will improve the presentation tremendously, and let one display one's application at resolutions of up to 800 X 600 with 64,000 colours. In the case of D-VR3 the application is set up to run on a 486/33 with VGA with Windows 95. The package will on run machines with Windows 3.1. mode but some aspects of the program are disabled in this mode.

   It is also advisable to work with a tool that supports Web deployment. Most multimedia tools will likely add Web access to their arsenals in the near future. Before choosing a tool for any long-term project, it is necessary to research the vendor's publicly announced plans for extending support to the Web. For this reason Toolbook is ideal as it supports a Netscape plug-in called Asymetrix Neuron. Neuron allows a designer to recompile and execute existing Toolbook applications over the Internet. Therefore, a multimedia based application like D-VR3 can be viewed over the Internet once Neuron is configured correctly.

5.17. Virtual Environments.

In the preceding sections the creation of VEs, how they are accessed and the tools used to navigate in them have been discussed. In this section the way in which these VEs are stored will be addressed. The storage of information on virtual objects and on the Virtual World itself is a major part of the design of any VE. The primary items that are stored in the World Database are the objects that inhabit the world, scripts that describe actions of those objects and the lighting, program controls, and hardware devices supported.

There are a number of different ways in which the world information may be stored: a single file, a collection of files, or a database. During the course of this research all three of these storage systems have been encountered. Each systems deals with files and actions in different ways. The single file approach was the most straightforward and organised. This was the type of storage system that was used by Virtus Walkthrough. The next type of storage system used was that used by the Id and Raven engines. Both these engines used a database system in which all the different files used to create different aspects of the Virtual World were stored in the one database. This database in both the ID and Raven engines was compiled into one extremely large file. The third type of approach, the multiple files approach, was the one used by the 3D Realms engine. The multiple file method is one of the more common approaches for VR development packages. Each object has one or more files (geometry, scripts, etc.) and there is an overall 'world' file that causes the other files to be loaded. Some systems also include a configuration file that defines the hardware interface connections.

5.18. Summary.

In this chapter the processes used to create VEs have been discussed. The design process described the steps that a designer should take in order to design a VE. A prototype model of the Computer Section, D.I.T. Cathal Brugha St., has been designed and then an actual model of the Monastic Enclosure at Glendalough. These models were designed using an array of software systems varying from engines used in 3D games to modelling and visualisation packages to 3D graphics and animation packages. Finally, each of the VEs created was incorporated into a multimedia-based environment.

This chapter has described in detail the processes that were taken in designing each VE constructed during the course of this research. The following chapter describes the evaluation of each of these VEs using a number of different processes.

Chapter Six

Evaluation of The Virtual Environments

6.1. Introduction.

Evaluation is an integral part of any design process and in the development of D-VR3 it is no exception. The purpose of this chapter is to evaluate each of the Virtual Environments (VEs) created from a number of different perspectives. The evaluation process, illustrated in Figure 6.1., was divided into two main sections;

a) The Internal Evaluation.
b) The External Evaluation.

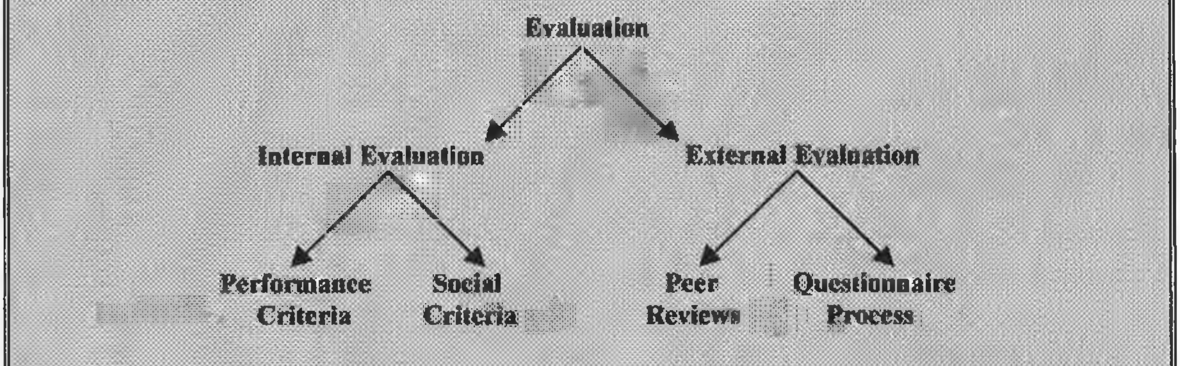

Figure 6.1. The Evaluation Process

The internal evaluation was a process which involved the evaluation of the specific features of the VEs constructed in chapter 5. The internal evaluation was conducted by drawing up a set of criteria against which each Virtual Environment (VE) was measured. The external evaluation, on the other hand, was a process involving the gathering of information from potential end-users. This information was gathered with the use of two methods:

a) Peer Reviews.
b) Questionnaire Process.

The first point that must be clarified when dealing with the evaluation process is deciding on exactly what is to be evaluated. In keeping with the definition chosen, section 1.2., the unit that was to be evaluated was to be the VE itself and not the processes used to create the environment nor the feeling of presence that a participant experiences within the environment.

6.2. Internal Evaluation.

The objective of the internal evaluation was to compare and contrast each VE in an attempt to determine whether they were compatible with each other and, if they were not, to choose the VE most suitable to the design of a tourism application. As an aid in this analytical process a strict set of criteria was drawn up. This set of criteria, illustrated in Table 6.1., was divided into two categories: Performance Criteria, which define the dimensions of VR, and Social Criteria, which describe how easily the application can be distributed and accepted into the market. Using this simple set of criteria all the VEs could be evaluated and placed on a 3D representation.

6.2.1. Performance Criteria.

It is not considered within the scope of this thesis to evaluate the social aspects of the VEs designed. The purpose of the evaluation process focuses on the Performance Criteria associated with each VE and the emphasis of the internal evaluation will, therefore, be focused on each environment's sensory vividness, interactivity and the whether it can be explored in real-time.

Table 6.1. Basic Criteria for Evaluating Virtual Environments.

Criterion	Description
(a) Performance Criteria.	
Sensory Vividness	The capabilities of the base hardware and software, and not just the output devices, contribute to sensory vividness, a key component in creating a high quality VE. Sensory vividness includes such considerations as (a) the number of sensory channels supported, and (b) the sensory resolution within each sensory channel.
Interactivity	Interactivity is also critical to create a high quality VE. In VR the criteria can be defined as (a) the number and forms of input and output, (b) the level of responsiveness to conscious and unconscious user actions and states, and (c) the range of interactive experiences offered by the system.
Time	Time refers to the time lag between the instruction and the resulting action; when there is no time lag in the system the system is known as performing actions in real-time. Real-time interaction is where the action of the user is immediately simulated by the mediated environment.
(b) Social Criteria.	
Sociability	Sociability is defined as the number of users a system can support. The lowest level of sociability is the single user interacting with the contents of a Virtual World. This is an important factor to consider when creating applications for the tourism industry as an applications sociability may greatly increase its usefulness.
Diffusability	Diffusability is the likelihood that a system can be adopted and used by various business, educational, and personal users. Variables that can raise or lower diffusability include: (a) Hardware compatibility: Systems that run on computers that are commonly available in organisations as opposed to those that require highly specialised computers. (b) Software compatibility: Systems that can import existing 2D and 3D models are more likely to be diffused. (c) Tool integration: Integrated systems with hardware support, modular software, and catalogues of virtual objects and worlds are more likely to be diffusable. (d) User friendliness: Systems that allow world building with minimal to moderate computing skills.
Cost	Cost is influenced by all of the variables above. One can anticipate the same pattern of development that has occurred elsewhere in the computer industry: cost dropping as low-end systems 'inherit' high-performance features previously available only on the most costly systems.

(Derived from Biocca 1992)

Each of the specific criteria used to evaluate the three types of performance criteria were weighted on a scale of 1 to 5 in accordance with their importance in improving a VE. In this scale 1 is awarded if a VE possesses a specific criterion which would make it a slightly better world and a weight of 5 would be awarded in the case of a criterion which is absolutely essential to a VE. The weightings for the evaluation criteria are displayed in Appendix J.

6.2.1.1. Sensory Vividness.

This is an extremely important factor to consider when evaluating a VE. This is the reason that these factors were weighted so heavily. Each of the packages evaluated supports both sound and vision except for Virtus Walkthrough and VRML Version 1.0. The criteria used to measure the sensory vividness of a VE resemble those used to create the actual VEs. Each criterion is weighted in order to reflect its importance to the VE. The first factor that must be considered when evaluating a VE is how many Sensory channels the world serves.

Number of Sensory Channels: The sensory channels are sight, sound, smell, taste and touch. In the case of each channel the Virtual Worlds were evaluated to display, not only whether the world actually serves this channel, but also at what resolution each channel was served. Each of these factors has a significant effect on the type of VE that is displayed and, therefore, they were weighted accordingly; this is discussed in Appendix J. None of the VEs evaluated served either the smell or taste channels as these have yet to be mastered. The touch channel was deliberately overlooked as the focus of this research is based on systems available on commercially available PCs with no additional hardware requirements. Therefore, the only two channels served by the VEs created were sight and sound. These channels were served to varying degrees by each of the worlds constructed. The worlds constructed with both the Id engine and the Raven engine served both of these channels with a medium resolution output while the worlds created using the 3D Realms engine and 3D Studio served both channels with a high resolution output and, finally, Virtus Walkthrough Version 2.6. Beta and VRML only serve the sight channel which they did with a high resolution output. Table 6.2. indicates the resolution that each sensory channel is served at in each VE. A low resolution was awarded one point, a medium resolution two marks and, finally, a high resolution was awarded three marks.

Table 6.2. Number of Sensory Channels.

	Id Engine	Raven Engine	3D Realms Engine	Virtus Walkthrough	VRML Version1	3D Studio
Sight	Medium	Medium	High	High	High	High
Sound	Medium	Medium	High	0	0	High
Smell	0	0	0	0	0	0
Taste	0	0	0	0	0	0
Touch	0	0	0	0	0	0
Total	20	20	30	15	15	30

Once the sensory channels had been evaluated the next criterion used was to examine how each VE tackled the geometry features. These features varied from the scaling involved in each world to whether each world used level of detail modelling. Again each of these features was weighted in accordance with their importance in VEs.

Geometry: The evaluation of the geometry features in any VE is a very important task as it is on this foundation that any good VE is constructed. For this reason, a lot of emphasis and, from an evaluation perspective, weight was placed on these specific criteria. The actual weight that each specific feature is attributed is shown in Appendix J. Two features in the geometry category have different levels included in their evaluation. They are scale and scaleable sprites. There were three different levels of scaling used to evaluate each VE; the first level was any scale which was not a recognised measurement system and could not be directly equated to any recognised system; the second level was any scale which was not a recognised measurement system but could be equated to a recognised system; and the third and final level was one which had a recognised measurement system. The first level was awarded a marking of one point, the second level a marking of two points and the third level a marking of three points. All of the VEs created had some form of scaling and those which were not recognised measurement systems could be equated to recognised measurements and, therefore, each world could be constructed to a precise scale. This criterion was weighted at 5 as it was deemed an essential part of a realistic Virtual World. There were also two levels of scaleable sprites used in the evaluation process; the first was a sprite that always faced the participant regardless of the participant's position in respect to the sprite. This type of sprite was awarded one point. The second type of scaleable sprite could be placed in a permanent position that would not change regardless of the position of the participant. This type of sprite was awarded two points. This criterion was weighted at 3 as it was deemed a useful addition to any world in regards to providing a participant with a sense of scale and filling the environment with furniture and characters.

Table 6.3. Geometry.

	Id Engine	Raven Engine	3D Realms Engine	Virtus Walkthrough	VRML Version1	3D Studio
Scale	Medium	Medium	Medium	High	High	High
Scaleable Sprites	Low	Low	High	No	No	No
Anti Aliasing	No	No	No	No	Yes	Yes
Level Of Detail	No	No	No	No	Yes	Yes
Animated Skies	No	Yes	Yes	No	Yes	Yes
Sector Over Sector	No	No	Yes	Yes	Yes	Yes
Slopes	No	No	Yes	Yes	Yes	Yes
Portals	Yes	Yes	Yes	No	Yes	No
Animated Furniture	Yes	Yes	Yes	No	No	Yes
Arches	No	No	No	No	No	Yes
Curved Surfaces	No	No	No	Yes	Yes	Yes
Total	20	22	34	28	42	45

The remainder of the geometric features did not have individual options included in them and were simply present within the VE or not. Each of these features were weighted in accordance with the weights displayed in Appendix J. The results which each VE received in regards to their geometric features are clearly indicated in 6.3.

Lighting: This section deals with the evaluation of the various techniques used to illuminate a VE. These techniques were previously discussed in section 5.3.2. These features vary from the essential ambient feature to the additional projector effect which only the worlds created using 3D Studio possessed. It was in this section that the environments created using 3D Studio really began to excel as they possessed all the lighting effects evaluated. The results that each VE received in this evaluation are clearly presented in Table 6.4.

Table 6.4. Lighting.

	Id Engine	Raven Engine	3D Realms Engine	Virtus Walkthrough	VRML Version1	3D Studio
Ambient	Yes	Yes	Yes	Yes	Yes	Yes
Directional	Yes	Yes	Yes	Yes	Yes	Yes
Multiple	Yes	Yes	Yes	Yes	Yes	Yes
Show Cone	No	No	No	No	No	Yes
Hot Spot	No	No	No	No	No	Yes
Shadows	Yes	Yes	Yes	No	No	Yes
Coloured Lighting	No	No	No	Yes	No	Yes
Projector Effect	No	No	No	No	No	Yes
Total	16	16	16	16	13	26

Shading: This section of the evaluation process deals with four different levels of shading and certain features associated with shading. The four levels of shading used to evaluate each VE were: *Wireframe*, which is effectively no shading; *Flat shading* where each object's face is filled in completely using the same colour; *Gouraud shading* which displays each face as a colour gradient using the surface properties at the vertex, its location, and its orientation relative to the lights; and, finally, *Phong shading* which calculates the color of each pixel using the properties at the pixel, its location, and its orientation relative to the lights. The three other features that were used as evaluation criteria were whether each environment could support reflective, refractive and transparent materials. The results of the evaluation of the shading effects used in each VE created are clearly displayed in Table 6.5.

Table 6.5. Shading.

	Id Engine	Raven Engine	3D Realms Engine	Virtus Walkthrough	VRML Version1	3D Studio
Wireframe	Yes	Yes	Yes	Yes	Yes	Yes
Flat	Yes	Yes	Yes	Yes	Yes	Yes
Gouraud	No	No	No	Yes	Yes	Yes
Phong	No	No	No	No	No	Yes
Reflective (Chrome)	No	No	No	No	No	Yes
Refractive (Glass)	No	No	Yes	Yes	Yes	Yes
Transparent (Coloured Glass)	No	Yes	Yes	Yes	Yes	Yes
Total	10	13	16	20	20	27

All VEs evaluated could display Wireframe and Flat shading while only worlds created using Virtus Walkthrough, VRML and 3D Studio could display Gouraud shading and only 3D Studio could display Phong shading. Since Phong shading is required for the creation of reflective materials, it was, therefore, not surprising that only 3D Studio could display reflective materials. All the worlds created could display refractive materials and with the exception of those created using the Id engine could display transparent materials. VEs created using 3D Studio were the only environments that could cope with all these features.

Textures: This section was a very diverse section of the evaluation process ranging from whether an environment could display simple plain textures through to whether it could support real life imported animations. All the Virtual Worlds evaluated measured up extremely well to each of the features evaluated except in the case of animated textures where worlds created using Virtus Walkthrough were the only worlds not able to display these and in the case of bump mapping where environments created using 3D Studio were the only worlds able to display this feature. Bump mapping uses the intensity of the bitmap image to affect the apparent smoothness of the material's surface. The results of this section are clearly displayed in Table 6.6.

Table 6.6. Textures.						
	Id Engine	Raven Engine	3D Realms Engine	Virtus Walkthrough	VRML Version1	3D Studio
Plain	Yes	Yes	Yes	Yes	Yes	Yes
Library	Yes	Yes	Yes	Yes	Yes	Yes
Real-life Import	Yes	Yes	Yes	Yes	Yes	Yes
Transparent	Yes	Yes	Yes	Yes	Yes	Yes
Alignment	Yes	Yes	Yes	Yes	Yes	Yes
Shading	Yes	Yes	Yes	Yes	Yes	Yes
Animated Textures	Yes	Yes	Yes	No	Yes	Yes
Mapping	Yes	Yes	Yes	Yes	Yes	Yes
Bump Map	No	No	No	No	No	Yes
Total	32	32	32	29	32	36

After evaluating each VE created against the above set of criteria it was then possible to group all these criteria together in order to determine the sensory vividness of each of the worlds created. Each VE's sensory vividness is clearly indicated in Table 6.7. The worlds created using 3D Studio were the most realistic in regards to sensory vividness with a rating of 74%, the worlds created using the 3D Realms engine were the next realistic with 58%, then those with VRML with 55%, then Virtus Walkthrough with 49% and finally, the Raven engine and the Id engine with 46% and 44% respectively. This rating while significant on its own must be used in conjunction with the ratings from Interactivity and Time in order for them to be viewed in the correct perspective.

Table 6.7. Sensory Vividness.						
	Id Engine	Raven Engine	3D Realms Engine	Virtus Walkthrough	VRML Version 1	3D Studio
Points Obtained	98	103	128	108	122	164
Total Points	222	222	222	222	222	222
Percentage	44%	46%	58%	49%	55%	74%

6.2.1.2. Interactivity.

In this part of the internal evaluation the specific criteria associated with the Interactivity of the VEs created is analysed.

Input Devices: One of the factors of Interactivity that must be considered is how flexible each VE is to adapt to the use of different input devices. Each VE was tested to determine whether it could accept a wide range of input devices from the conventional mouse to thought transducers. The more common input devices suited to commercially available PCs were weighted more heavily than the more complex devices. Table 6.8. indicates which input device each world supports. The worlds created using the 3D Realms engine accepted the most input devices. These worlds accepted nine out of the fourteen input devices identified. Worlds created using 3D Studio, Virtus Walkthrough and all accepted eight devices and both the Id engine and the Raven engine accepted six devices. None of the VEs evaluated accepted the more complex input devices such as data gloves, bodysuits, retinal scanners, speech transducers or thought transducers. This was to be expected as these devices are not aimed at the desktop VR market and, therefore, many desktop VR engines would not be equipped to cope with such devices.

Table 6.8. Input Devices - Transducers.

	Id Engine	Raven Engine	3D Realms Engine	Virtus Walkthrough	VRML Version1	3D Studio
Mouse 2D	Yes	Yes	Yes	Yes	Yes	Yes
Keyboard	Yes	Yes	Yes	Yes	Yes	Yes
Joystick	Yes	Yes	Yes	Yes	Yes	Yes
Mouse 3D	Yes	Yes	Yes	Yes	Yes	Yes
Trackball	No	No	Yes	Yes	Yes	Yes
Spaceball	No	No	Yes	Yes	Yes	Yes
Touchscreen	Yes	Yes	Yes	Yes	Yes	Yes
Wand	No	No	Yes	Yes	Yes	Yes
Sensors	Yes	Yes	Yes	No	No	No
Gloves	No	No	No	No	No	No
Bodysuits	No	No	No	No	No	No
Retinal Scanners	No	No	No	No	No	No
Speech Transducers	No	No	No	No	No	No
Thought Transducers	No	No	No	No	No	No
Total	24	24	33	30	30	30

Table 6.9. Output Devices - Displays.

	Id Engine	Raven Engine	3D Realms Engine	Virtus Walkthrough	VRML Version1	3D Studio
Computer Monitor	Yes	Yes	Yes	Yes	Yes	Yes
Head Mounted Display	Yes	Yes	Yes	Yes	No	Yes
BOOM	No	No	No	No	No	No
Red/Blue Glasses	No	No	Yes	Yes	No	No
Shutter Glasses	No	No	Yes	Yes	No	No
Cyberscope	Yes	Yes	Yes	Yes	Yes	Yes
Projected Image	Yes	Yes	Yes	Yes	Yes	Yes
Sound Device	Yes	Yes	Yes	No	No	Yes
Stereo Sound Display	Yes	Yes	Yes	No	No	Yes
Haptic Displays	No	No	No	No	No	No
Modems	Yes	Yes	Yes	No	Yes	No
Network	Yes	Yes	Yes	No	Yes	No
Serial/Parallel	Yes	Yes	Yes	No	Yes	No
Total	32	32	38	21	20	23

Output Devices: The interactivity feature of any VE is also dependent on the range of output displays that the world can support. Each of the VEs was evaluated against a number of output displays ranging from the conventional computer monitor to haptic displays in order to determine the range of output devices that it supported.

Of all the VEs evaluated those created using the 3D Realms engine supported the most output displays, eleven out of a possible thirteen devices. Worlds created using the Id engine and the Raven engine both supported nine devices; those created Virtus Walkthrough, VRML and 3D Studio all supported six devices out of the thirteen devices evaluated. Each device was individually weighted in order to reflect its importance to the VEs. Table 6.9. indicates which output devices each VE supports.

Level of Interactivity: The level of interactivity available to a participant in a VE is an extremely important factor to consider when evaluating the interactivity. The level of interactivity measures the amount of control the participant has to change the Virtual World. There are, in all, three levels of interactivity. The first level involves a participant who may look at, but not touch, objects within a Virtual World. The second level of interactivity is when the participant gains, or is granted, more power which enables him to modify entities or processes in the environment, either temporarily or permanently, but he may not change the base world. The third and highest level, of interactivity, is where the individual is granted the ability to create new worlds including the base environment. Table 6.10. indicates each environment's level of interactivity.

Table 6.10. Level of Interactivity.

	Id Engine	Raven Engine	3D Realms Engine	Virtus Walkthrough	VRML Version1	3D Studio
Look at, but not touch, objects	Yes	Yes	Yes	Yes	Yes	Yes
Interact with objects	Yes	Yes	Yes	Yes	Yes	No
Interact with and change objects	No	No	No	Yes	No	No
Total	10	10	10	15	10	5

Degrees of Freedom: Not only is it important to evaluate each VE to determine which devices it supports but it is also important to determine the movement of a participant in a VE in terms of his position and orientation. In other words, it is important to determine whether a VE is capable of measuring a participant's position in terms of the X, Y and Z axis and his orientation in terms of pitch, roll and yaw. This is alternatively known as a VE's degrees of freedom. This is another important feature to consider when determining the interactivity and, indeed, the overall realism of any Virtual World. For this reason each of these criteria is weighted at 5. Worlds created using Virtus Walkthrough, VRML and 3D Studio all have 6 DOF, worlds created with both the Raven engine and the 3D Realms engine both have 5 DOF, and finally, the worlds created using Id engine have 3 DOF. The results of this evaluation are clearly illustrated in Table 6.11.

Table 6.11. Degrees of Freedom.

	Id Engine	Raven Engine	3D Realms Engine	Virtus Walkthrough	VRML Version1	3D Studio
Left/ Right (x)	Yes	Yes	Yes	Yes	Yes	Yes
Forward/Backward (z)	Yes	Yes	Yes	Yes	Yes	Yes
Up/Down (y)	No	Yes	Yes	Yes	Yes	Yes
Pitch	No	Yes	Yes	Yes	Yes	Yes
Roll	No	No	No	Yes	Yes	Yes
Yaw	Yes	Yes	Yes	Yes	Yes	Yes
Total	15	25	25	30	30	30

Specific Interactivity Features: This section evaluates each VE created against a number of individual interactivity effects. These effects range from whether a sliding door could be used within a VE to whether moving sectors could be used within a VE. Worlds created using the 3D Realms engine were the only worlds capable of producing all the effects evaluated. Environments created using the Raven engine coped moderately well with these effects and worlds created using the Id engine could only produce sliding doors. Worlds created using any of the other three engines could not produce any of the effects evaluated. These results are presented in Table 6.12.

Table 6.12. Specific Interactivity Features.

	Id Engine	Raven Engine	3D Realms Engine	Virtus Walkthrough	VRML Version1	3D Studio
Sliding Doors	Yes	Yes	Yes	No	No	No
Swinging Doors	No	Yes	Yes	No	No	No
Mirrors	No	No	Yes	No	No	No
Moving Sectors	No	Yes	Yes	No	No	No
Earth Movements	No	Yes	Yes	No	No	No
Demo Cameras	No	No	Yes	No	No	No
Breakable Glass	No	Yes	Yes	No	No	No
Echo	No	No	Yes	No	No	No
Total	4	17	26	0	0	0

Table 6.13. Scripting.

	Id Engine	Raven Engine	3D Realms Engine	Virtus Walkthrough	VRML Version1	3D Studio
Motion Scripts	Yes	Yes	Yes	No	No	No
Trigger Scripts	Yes	Yes	Yes	Yes	Yes	No
Connection Scripts	No	No	No	No	No	No
Procedural Modelling	No	No	No	No	No	No
Simple Animation	Yes	Yes	Yes	No	Yes	Yes
Interaction Feedback	Yes	Yes	Yes	Yes	Yes	No
Total	16	16	16	8	12	4

Scripting: The scripting section evaluates each VEs in terms of the scripts it supports. All the different types of scripting, previously discussed in section 5.4., were weighted at 4 as none of them are an essential part of a VE. As is indicated in Table 6.13. the Id engine, the Raven engine and the 3D Realms engine all support four scripting types, VRML supports three scripting types, Virtus Walkthrough supports two types and, finally, 3D Studio only supports one scripting type.

After evaluating each VE created against the interactivity criteria it was then possible to group all of these criteria together in order to determine each world's interactivity. Each VE's interactivity is clearly indicated in Table 6.14. The worlds created using the 3D Realms were by far the most interactive with a rating of 80%, the worlds created using the Raven engine were next with 67%, then those created with Virtus Walkthrough with 56%, then VRML with 55% and finally, the Id engine and the 3D Studio with 55% and 50% respectively. These results again must be analysed in conjunction with the results from the other Performance criteria in order for them to be viewed in the correct perspective.

Table 6.14. Interactivity.						
	Id Engine	Raven Engine	3D Realms Engine	Virtus Walkthrough	VRML Version 1	3D Studio
Points Obtained	101	124	148	104	102	92
Total Points	185	185	185	185	185	185
Percentage	55%	67%	80%	56%	55%	50%

6.2.1.3. Time.

The third, and final, Performance Criteria is time. In a VR system, lag is of vital importance since it measures the time taken between a movement being initiated at the input device, and the resulting action being displayed to the participant. The lag of a VE can be calculated by the summation of its two component parts *Latency* and *Update Rate*.

Table 6.15. Latency.						
	Id Engine	Raven Engine	3D Realms Engine	Virtus Walkthrough	VRML Version1	3D Studio
Latency	5	5	5	5	5	5
Total	25	25	25	25	25	25

Latency: Latency is the rate at which the acquisition portion of the system can acquire new data (Earnshaw 1995). As such it deals with signals in the real world and is an evaluation of the input process to perform in real-time. There are two aspects to this: how accurate the input device is, and secondly, how quickly the Virtual World can sense a change from the input device. The latency of each VE is clearly indicated in Table 6.15. In this evaluation points from 1 to 5 are awarded for each VE's latency. A score of 1 means that the world is very slow to pick up movement of the input device and the input device is not very accurate. A score of 5 means that the input device is very accurate and performs in real-time. In all cases the engines were very responsive to this criterion.

Update Rate: The Update Rate is a VE's ability to present position and orientation data to the participant. Many users feel that a high update rate is a very important factor in any VE. Update rate is not a function of the acquisition process and consequently does not reflect overall system performance (Kalawsky 1993). The update rate is normally measured in frames per second (fps) or in polygons per second (polys/sec). Fps refers to the number of frames appearing on the display each second - the speed of the VE. In a cinema a film runs at about 24 fps, on television, programmes run at about 30 fps, cartoons on television are displayed at 15 fps. In all of these, an individual can not distinguish individual frame; instead they move by in a sequence which is perceived as a continuous motion. Unlike films, programmes and cartoons, VEs exist in real-time. They are not played back from a recording medium. The desired frame rate in VEs is in the region of 22 fps (Jacobson 1994). Polys/sec is often quoted as a performance measure of VR software. The figure simply describes how many polygons could be processed in a second, and provides a broad idea of how many frames per second would be achievable for a given world. The Update Rate of each VE is indicated in Table 6.16. Since some of the VEs examined did not possess a mechanism for measuring either fps or polys/sec, the Update Rate of each had to be attributed by using subjective means. A score of 5 means that the world has a high update rate and 1 means that the world has an unacceptably low update rate.

Table 6.16. Update Rate.						
	Id Engine	Raven Engine	3D Realms Engine	Virtus Walkthrough	VRML Version1	3D Studio
Update Rate	4	4	5	3	3	5
Total	20	20	25	15	15	25

The evaluation of time in a VE was a difficult process. This was due to the different types of engines adopted to construct the VEs. Each VE's time criteria are indicated in Table 6.17. A VE which has no appreciable lag involved is said to perform in real-time. Virtual Worlds created using the 3D Realms engine and 3D Studio both perform in real-time, no phase lag, with the Raven engine and the Id engine both creating VEs which perform at 90% of real-time and, finally, worlds created using Virtus Walkthrough and VRML performing slightly slower again at 80% of real-time.

Table 6.17. Time.						
	Id Engine	Raven Engine	3D Realms Engine	Virtus Walkthrough	VRML Version1	3D Studio
Points Obtained	45	45	50	40	40	50
Total Points	50	50	50	50	50	50
Percentage	90%	90%	100%	80%	80%	100%

6.2.2. Conclusion to the Internal Evaluation.

In the conclusion of sections 6.2.1.1, 6.2.1.2. and 6.2.1.3. dealing with the evaluation of each VE against sensory vividness, interactivity and time respectively, it was stated that individual results had to be analysed in conjunction with all the other Performance Criteria. In this section the results that each VE obtained in each of the previous sections will be discussed in conjunction with all the other appropriate results. The overall results that each VE obtained are illustrated in Table 6.18. and Figure 6.2. These grades are presented in percentage form where 100% would be the ideal sensory vividness, ideal interactivity and ideal time of the VE. Furthermore, this method of evaluation is useful as a means of a comparative analysis between VEs. It can, therefore, be concluded that the VEs created using 3D Studio had the best sensory vividness at 74% of the ideal but had very poor interactivity, at 50%. VEs created using the 3D Realms engine dealt with all the criteria evaluated very well. The VE's achieved the second highest sensory vividness at 59% of the ideal, the highest interactivity evaluated at 80% of the ideal, and the joint highest time at 100% of the ideal.

Table 6.18. Overall Performance Criteria.						
	Id Engine	Raven Engine	3D Realms Engine	Virtus Walkthrough	VRML Version1	3D Studio
Sensory Vividness	44%	46%	58%	49%	55%	74%
Interactivity	55%	67%	80%	56%	55%	50%
Time	90%	90%	100%	80%	80%	100%

The VEs created using VRML were the only other worlds that provided an above average sensory vividness at 55%. These worlds, on the other hand, had a very low interactivity rating in comparison to the other VEs evaluated at 55% and had a less than average rating in the time criteria. The worlds created using Virtus Walkthrough scored a less than average rating on all three criteria. Worlds created using the remaining two games engines, the Id engine and the Raven engine, both obtained less than average ratings for sensory vividness and average ratings for time. VEs created using the Raven engine had very good interactivity, 67% of the ideal while those created using the Id engine had a less than average Interactivity at 55%. The results that each VE obtained were placed on a 3D representation, Figure 6.2., in order to illustrate how each world performed in comparison to the other worlds evaluated.

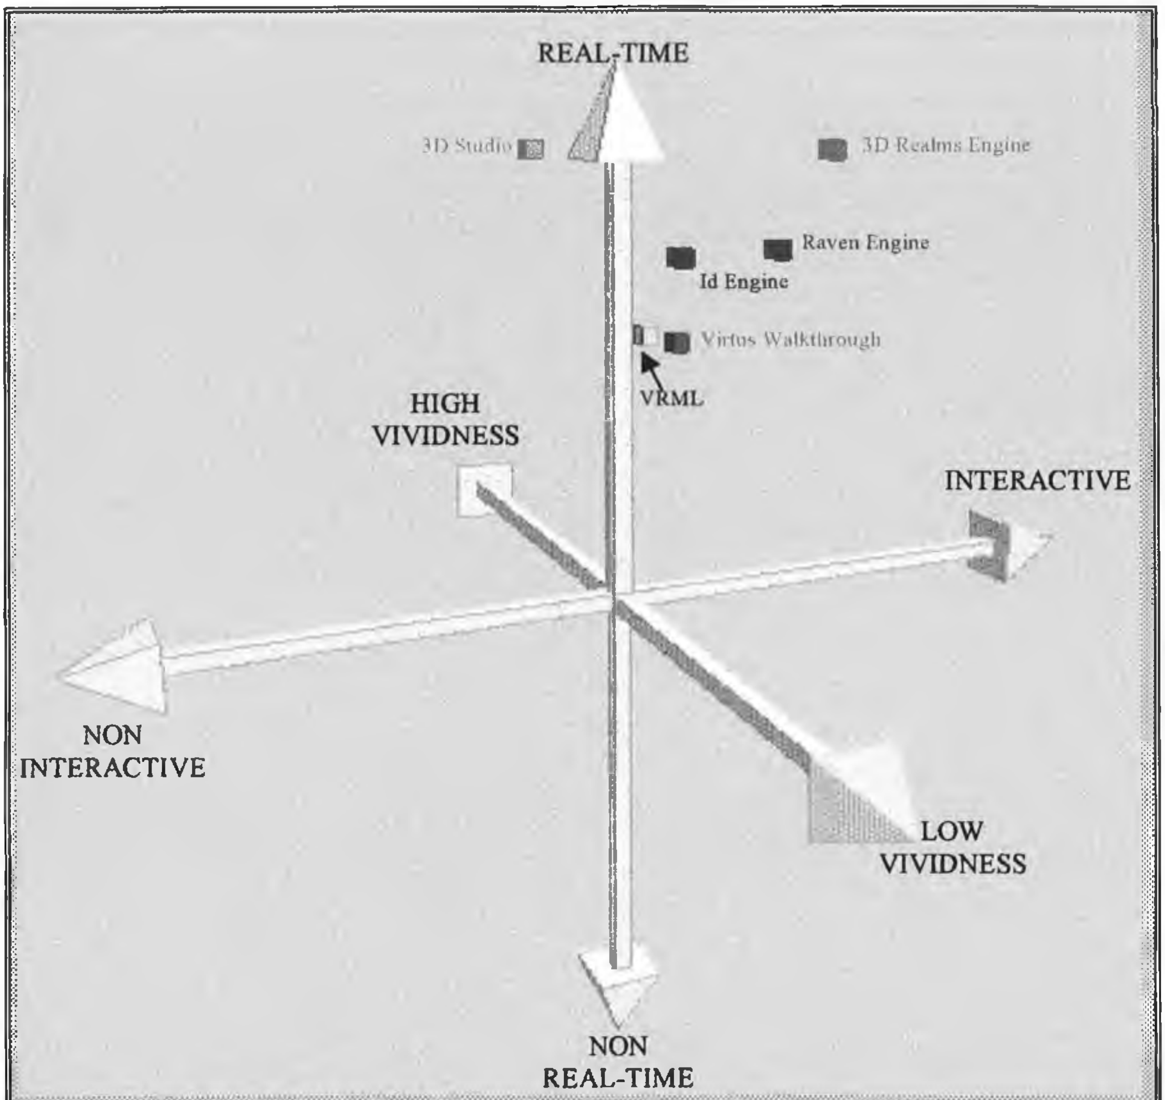

Figure 6.2. The Virtual Environments Plotted on a Three Dimensional Representation.

Each VE became more useful as a tourism decision-making tool as the sensory vividness and the interactivity of the model increased. However, attention must also be paid to the effect that both of these criteria have on the speed of a VE. In real-time interactive VEs a trade-off exists between the complexity of the VE and the performance of the VE. The greater the complexity of the VE, the slower the frame rate. Cruder VEs mean a better frame rate and less lag involved.

The VEs created using the three games engines all put the burden of rendering on the participant's machine, the Image generator, which means lower quality, blockier images without real reflections or refraction. However, in these environments the participant has much greater freedom of movement because the computer calculates in real-time each new viewpoint the participant sees. The three games engines that were evaluated feature interactive graphics which store only the comparatively small shapes and textures of the Virtual World and, thus, take up only a small number of megabytes of storage space. This is in sharp contrast to the hundreds of megabytes needed to store fully rendered frames of an animation such as those produced by 3D Studio.

In an attempt to provide high interactivity on commercially available PC's, the developer typically has to give up some desirable features. For instance, the low-resolution textures used on walls are often very small and lack detail, leaving the image looking blocky. In exchange, the developer gets increased animation speed and some assurance that image-redraw speed will not interfere with his movement through the world.

One can conclude this section by stating that vividness, interactivity and time play against one another. This trade-off between these three criteria is made to some degree in all current VR systems in that a reduction in the need for interactivity allows for a higher level of realism. A solution would be to render only what would be in the view of the participant and, therefore, the toll on the system would be reduced and the system therefore, speeds up and the effect of realism is enhanced.

6.3. External Evaluation.

This section presents an analysis of the different types of research method in an attempt to determine which is the most appropriate for use within this thesis.

6.3.1. Examination of Research Instruments.

Research instruments are the methods that can be used to collect the information that is required by the researcher. The different methods of primary research, outlined in Table 6.19., are reactive methods, in which the researcher may intrude on the environment, and non-reactive methods, in which the researcher has no-interaction with the environment. It was decided that the methods of investigation and information collection that best suited the information required in this thesis were to be peer reviews and questionnaire survey.

Table 6.19. The Methods of Primary Research.

Reactive Methods	**Non-Reactive Methods.**
1. Participant Observations.	1. Archives.
2. Focus Group Discussions.	2. Observation.
3. Surveys.	3. Physical Traces.
4. Experiments.	

(Kotler 1991)

6.3.2. Peer Reviews.

The main reasons for using peer reviews were to ensure that the application that was developed was of a usable quality for end-users, to encourage respondents to discuss topics in detail rather than to just answer questions put to them, and finally, to act as an aid for designing an effective questionnaire. The first step taken regarding the peer reviews was to provide the members of each group with the relevant information about D-VR3 before in-depth discussions about this topic could begin. In order to achieve this an audio-visual presentation accompanied by a demonstration of the actual application was used.

The composition of the peer reviews has to be related to the aims and objectives of the research. The peer review process was approached from two different perspectives. The first peer review was conducted with our Industrial Sponsors, Siemens Nixdorf. This peer review was held on Tuesday the 8th of August 1996 in the Computer Section, Cathal Brugha St. The group was comprised of Mr. Brendan Kelly, Sales Manager, Siemens Nixdorf, Mr. Stephen Loughney, Sales Executive Siemens Nixdorf, Dr. Ciaran MacDonaill and the author. The main aim of this peer review was to present the application to our sponsors and to gain as much feedback, preferably of a technical nature, as possible. The second peer review held on Friday the 23rd of August 1996, also took place in the Computer Section, Cathal Brugha St.. On this occasion the group members were comprised of people from a tourism background and were Mr. George McClaferty, manager of the interpretive center in Glendalough, Ms. Claire Tuffy, manager of the Interpretive Center at Newgrange, Mr. Willy Cummins, Architect for the Office of Public Works (OPW), Ms. Catherine Ward also from the OPW and the author. The main aim of this presentation was to gain as much feedback about the application, this time from a tourism perspective. In an attempt to acquire as much information from each peer review as possible, the discussions lasted in excess of one and a half hours each.

On both occasions the application was well received and some extremely valuable lessons were learned. The main finding which arose from the peer reviews were:

1. A more user-friendly, intuitive input device was required. The reason was that the mouse and keyboard approach, the input devices which had been employed up until this point, were not intuitive enough for an absolute beginner to comprehend and command successfully. For this reason a joystick was employed and was found to be more user-friendly and intuitive.
2. The use of background music was unacceptable. For this reason, the use of background music was removed from D-VR3.
3. More realistic textures were required. Until this point the textures used were the actual scanned textures from each building but only one texture has been applied to the whole building. It was felt that this process was unacceptable and that a more realistic texture application process was required. This process was examined, a more acceptable process, entailing the application of several textures from each part of the buildings to the models, was applied to each of the models constructed.
4. It was revealed that it would be beneficial to some parties if one could see the development of each building through the centuries. While this would indeed be interesting it was not considered a priority in this thesis.
5. It was also suggested that the walls which divided each of the models in Glendalough should be removed as they were somewhat misleading to anyone who viewed the models. Models of this nature had already been constructed but unfortunately had to be discarded because of the unacceptable update rate of some of the engines examined. This was due to the fact that each engine had to redraw

the entire Monastic Enclosure each time the participant moved within the VE whereas in the case of each building which is surrounded by walls the engine only has to redraw one model when the participant moves.

6. Another suggestion was that some objects should be added to the VE to provide the participant with a better sense of scale. In order to achieve a better sense of scale computer-operated characters and trees illustrated in Figure 5.40. were employed. This unfortunately was only possible within some of the engines evaluated.

These proposals were implemented prior to the questionnaire process . Group discussions yield results of a qualitative nature, not a quantitative nature. If quantitative results are required one must look to another form of research type. The reason for this is two fold - the samples are inevitably small and the actual form of the groups did not lend themselves to precise quantification of data as the discussions which took place were difficult to translate into measurable responses.

6.3.3. Questionnaire Process.

The questionnaire process entails formulating a basic plan which specifies the type of information to be collected, the sources of this information, the questionnaire design, the population and sample plan that were to be used in the research. An effective questionnaire process will ensure that the relevant information is collected by accurate and economical procedures. In this section the different stages of the questionnaire process are discussed.

6.3.3.1. Questionnaire Design.

For virtually every conceivable question there are several possible, and theoretically acceptable, forms; in choosing between them, knowledge of the questionnaire population, common sense, past experience and pilot work are at present the designer's main tools. The first step in designing a questionnaire was to decide on what information was required. Once this was achieved it was then possible to design the specific questions needed to obtain this information. It was essential in this case to keep the questionnaire length to a minimum because the respondents would all be tourists with time restrictions.

6.3.3.2. Question Types.

For the purpose of gaining the maximum amount of information from the questionnaire four types of question were used:

1. **Dichotomous.** Dichotomous questions are questions which require a straight "Yes" or "No" answer. These questions are easy to ask, understand, record and analyse. However, they do not yield detailed information and exclude "Don't Knows" and people without a defined "Yes/No" answer. These questions are very useful in obtaining factual information but they are of limited value when seeking opinions or attitudes; hence, the use of multiple choice and open-ended questions.

2. **Multiple Choice.** Multiple choice questions offer the respondents a number of alternatives and so permit the collection of more detailed data. As a number of alternatives are listed, it eliminates the problem of memory associated with the questionnaire and influences the respondent's thoughts in the direction desired. Response to these questions also facilitates cross analysis.

3. **Open-ended.** Open-ended questions, also referred to as attitudinal questions, give the respondent complete freedom in answering and, thus, are expected to yield the maximum information, although this is not always the case. They are used when it is impossible, or impractical, to formulate all the alternatives and where the respondent's opinion is sought. The main limitation of open-ended questions is that the respondent may not fully understand the information that is required, which may result in a low response to the questions or in the provision of information which is not particularly relevant to the area discussed.

4. **Polymorphous.** Polymorphous questions are questions which display characteristics of two or more of the above types.

The questionnaire, Appendix K, consists of eighteen short questions, some of which contain both filter questions requiring "Yes" or "No" answers and sub-sections which need or need not, be answered depending on the answer to the filter question. Appendix L outlines the nature and purpose of each question that appears in the questionnaire.

6.3.3.3. Questionnaire Population.

For the purpose of this thesis the questionnaire population is best described as being

"any visitor to the Interpretive Centre in Glendalough."

The number of visitors to the Interpretive Centre in 1995 was 101,779 visitors (Office of Public Works 1996), and it was considered impractical to study the complete population of interest and, therefore, a sample was required.

6.3.3.4. Questionnaire Sample.

Sampling is a method of gaining a representative experience of a subject area. A sample is a representative group within a population that reflects with reasonable accuracy the opinions, attitudes or behaviour of the entire group. In the case of this thesis the population was so large that to study everyone would be impractical. However, a sample of 216 respondents was investigated. Sampling is based on two fundamental principles of statistical theory.

1. The Law of Statistical Regularity - Any subject of a population selected at random will tend to hold the same characteristics as those of a larger group.

2. Law of Inertia - Larger groups are more stable than smaller groups, owing to the compensation effect or deviation in opposite directions.

However, taking a sample is not simply a matter of taking the nearest item. If worthwhile conclusions relating to the whole population are to be drawn from the results, it is essential to ensure as far as possible that the sample is free from bias. Bias is the influence in the selection of a sample of a particular feature in excess of its true importance. The possibility of taking a sample being biased can be reduced by taking a random sample. Given the nature of the population, Systematic sampling where the one in K^{th} sampling units was employed was thought to be the most appropriate.

6.3.3.5. Pilot of Questionnaire.

Questionnaires do not emerge fully fledged, they have to be created or adopted, fashioned and developed to maturity after test flight (Oppenheim 1992). For this reason a questionnaire must be tested prior to the actual survey to ensure that the correct information is being gathered by the questionnaire. The process of designing and testing questions and procedures is referred to as pilot work (Oppenheim 1992). The questionnaire involved in this thesis was piloted on a sample group of 15 respondents. The pilot questionnaires were administered in a similar medium to that which was to be used for the final questionnaire. As a result of these pilot tests it was believed that the questionnaire was ready for administration.

6.3.3.6. Questionnaire Administration.

The questionnaires were administered at the Interpretative Center at Glendalough between Wednesday 25th September 1996 and Friday 27th September 1996. The questionnaires were administered in conjunction with a short presentation of D-VR3 and its features - tourists were allowed to use the applications and many people had questions which were answered fully. Again, the application was well received and generated much interest over that period of time. In all 216 questionnaires were completed successfully.

6.3.3.7. Method of Analysis.

The results obtained from the questionnaire process were analysed by the author with the aid of an analytical computer package entitled Statistical Package for Social Sciences (SPSS). This package facilitated the production of frequency tables, cross tabulations and graphs, by inputting the data obtained from each questionnaire into the package.

6.3.3.8. Results and Analysis.

Question 1 – Do you currently have access to a computer? This question was designed in order to determine the number of respondents who have access to a computer. This information was required to determine how these respondents could be targeted if D-VR3, or a similar application, was commercialised. This question is a filter question and, therefore, only those who responded positively to this question were requested to answer the next question.

It was found that 76% of the respondents had access to a computer. These findings, illustrated in Figure 6.3., were encouraging as they showed that over three quarters of the sample analysed had access to a computer at one location or another and these respondents would, therefore, theoretically have access to D-VR3, or a similar application, were it commercialised. The findings from this question were further analysed by cross tabulating them against the information collected from the questions on gender, age group, tourist type, occupation and nationality. After these cross tabulation had been analysed no significant trend was detected.

Figure 6.3. How Many Respondents Have Access To A Computer.

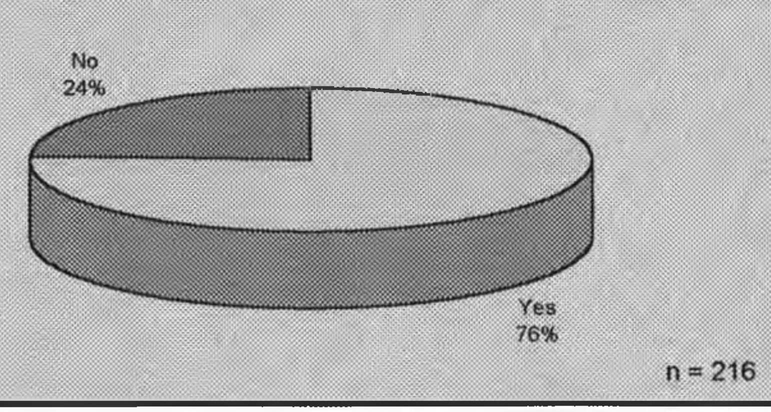

Question 2 – Where do you have access to a computer? As mentioned in the previous section, this is a question that is only to be answered by those who responded positively to Question 1. The aim of this question was to determine where exactly each respondent had access to a computer. This question was divided into four sections: those who had access at home, at work, at college and those who had access elsewhere. Those who responded that they had access to a computer elsewhere were prompted to supply where they had such access by use of an open-ended question.

The findings, illustrated in Figure 6.4., were enlightening as they showed that 71% of the respondents had access to a computer at home. This was a high percentage compared to only 49% of the respondents who had access at work and 18% of the respondents who had access to a computer at college. The remaining 1% of the respondents had access to a computer at a local library. This information would be of great importance if D-VR3, or a similar application, was to be commercialised. These results were also cross-tabulated against gender, age groups, tourist types, occupation and nationality and again no significant trend was detected.

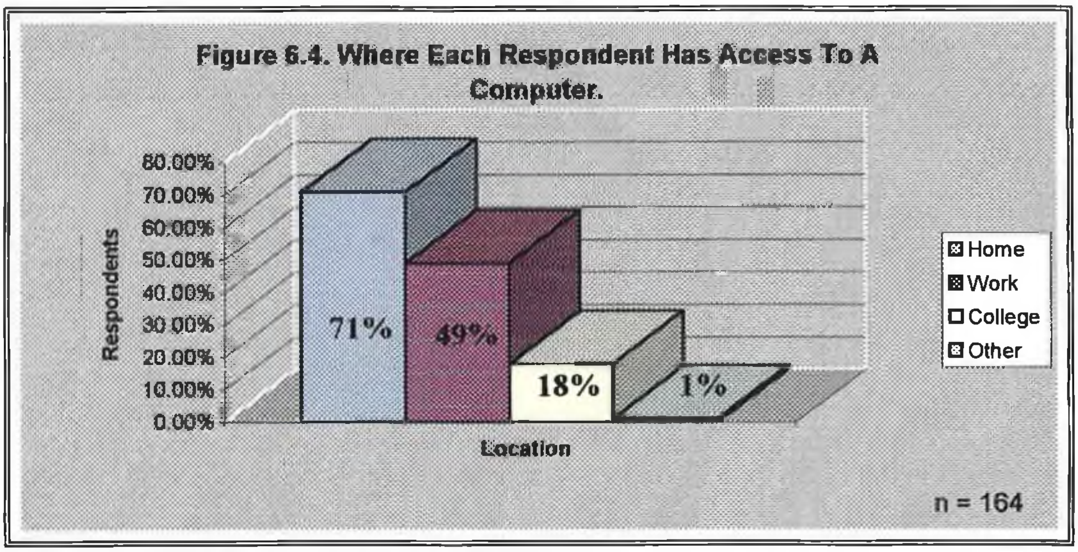

Question 3 – Would you be likely to use this program as a form of tourism information? This question was designed to determine how many of the respondents, after being presented with a demonstration of the functions and features of D-VR3, would use it in the future as a source of tourist information. This question is also a filter question as only those who respond positively to this question are requested to answer the following three questions.

This question was answered by all 216 respondents. A total of 97% of the respondents would use D-VR3 as a form of tourist information in the future. The results of this question are clearly illustrated in Figure 6.5. This extremely high percentage reflected the enthusiasm that D-VR3 generated over the three-day questionnaire administration and is a good measure of the application's popularity among the respondents. The results from this question were also cross-tabulated against the results obtained from the gender, age group, tourist type, occupation and nationality questions. Due to the extremely small negative response rate it was very difficult to spot a trend in any of the cross tabulations undertaken.

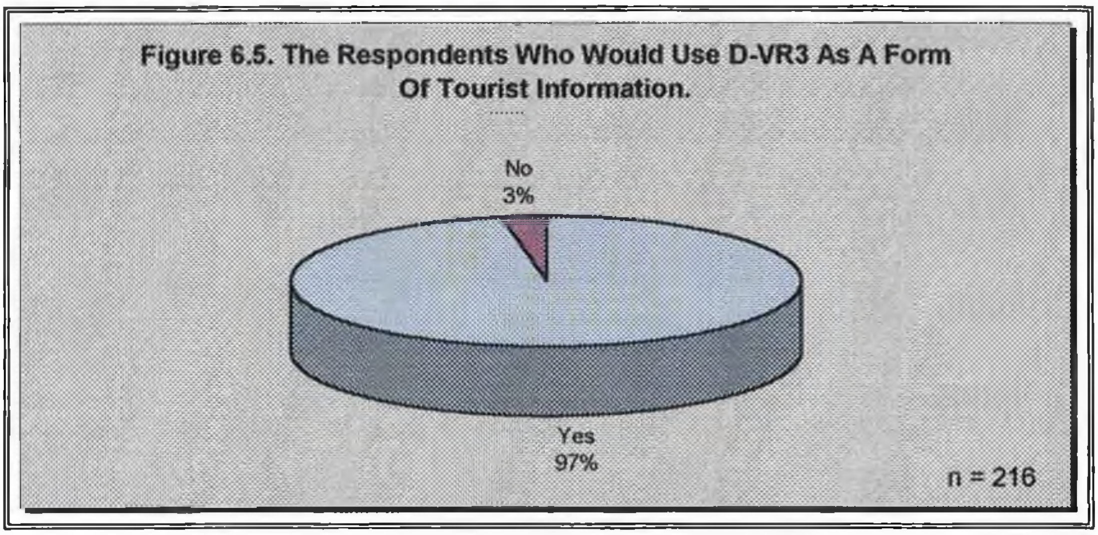

Question 4 – When would you be likely to use this program as a form of tourism information? The second part of this filter question was to determine whether those who would use D-VR3 as a form of tourist information in the future would they do so prior to their visit, during their visit or after their visit.

This question allowed each respondent to choose one, two or all of three options. The results obtained from this question are graphically displayed in Figure 6.6.

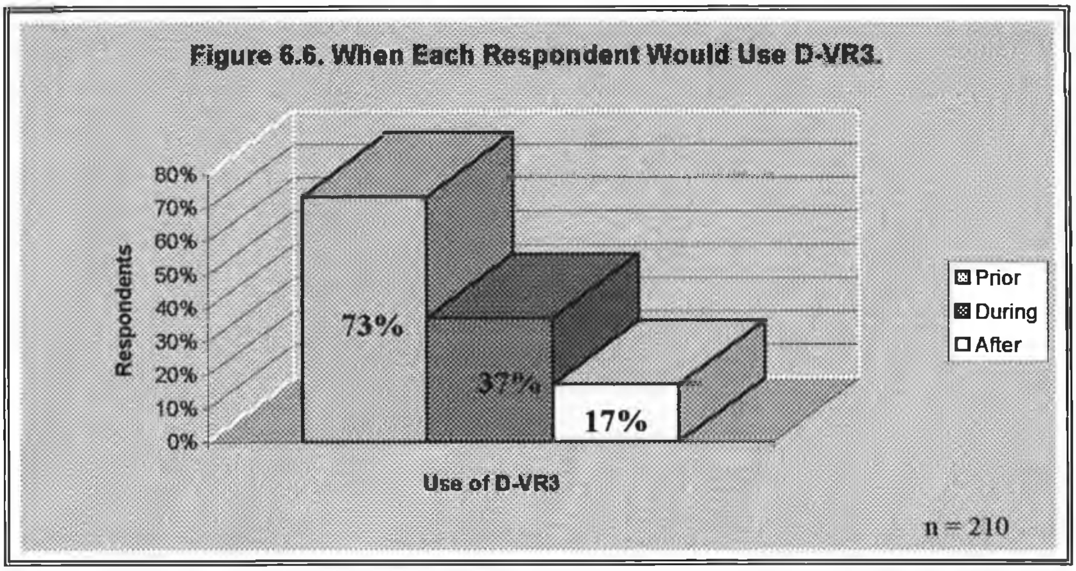

Figure 6.6. When Each Respondent Would Use D-VR3.

Of the 210 respondents who would avail of D-VR3 as a source of tourist information in the future should the opportunity arise, 73% of these would use it prior to their visit, 37% would use D-VR3 during their visit and, finally, 17% would use the application after their visit. It is clearly visible from Figure 6.6. that the respondents who would use D-VR3 prior to their visit presumably as a form of decision making and planning tool by far outnumber those who would use it during and after their visit. From this analysis one can conclude that the biggest use of D-VR3 as a form of tourist information should be targeted at tourists prior to them going to a tourism location.

Question 5 – Why would you be likely to use this program as a form of tourism information? At this stage in the analysis it has been established that 97% of the respondents said that they would use D-VR3 again. These respondents have subsequently been asked when they would use D-VR3. The next logical step was to determine why each of these respondents would opt to use this type of application as a source of tourist information. This question was the third part of the filter question in that only those people who responded positively to Question 3 were requested to answer this question. An open-ended question was used in order to allow each respondent enough depth to properly answer the question.

The reasons that each respondent gave for using D-VR3 as a source of tourist information are clearly indicated in Figure 6.7. A total of 32% of the respondents did not respond to this question at all. Of those that did respond to this question 36% would use D-VR3 because they believe that it provides more information than traditional tourist information sources, 27% would use D-VR3 because they believe that it provides a good overall impression of the featured tourist site and 19% would use D-VR3 because they believe that it would help them to plan their visit once they are at the site. A total of 5% of the respondents would use D-VR3 because they believe that the interactivity that it provides would make the whole experience of planning a visit to any destination more interesting. A total of 3% of the respondents would use D-VR3 because it was an interesting concept, 3% believed it was useful as a decision making tool, and another 3% thought would be very useful to use D-VR3 in the case of bad weather. Finally, 1% of the respondents would use D-VR3 because it provided a user with very good access to each of the buildings modelled, 1% would use it because they believed it to be a very useful learning tool about the tourist sites, and, finally, another 1% would use the application in order to view the buildings as they originally appeared.

Figure 6.7. Why Each Respondent Would Use D-VR3.

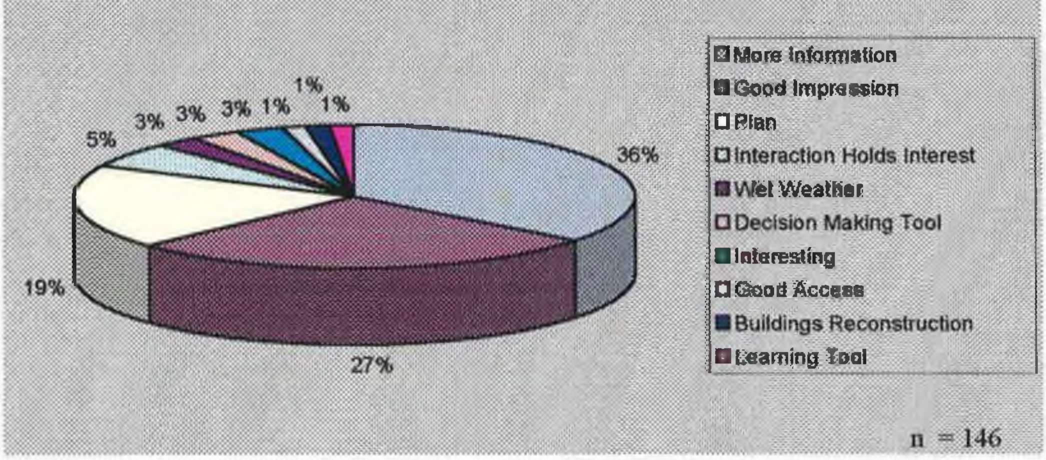

Question 6 – Where would you be likely to use this program as a form of tourism information? This question is the final part of the filter question that began with Question 3. The aim of this question was to establish where each respondent would be like to use D-VR3. This question allowed each respondent to choose one location, or a variety of locations, in which they would use such a program. The different locations that people chose are clearly indicated in Figure 6.8.

Figure 6.8. Where Each Respondent Would Be Likely to Use D-VR3.

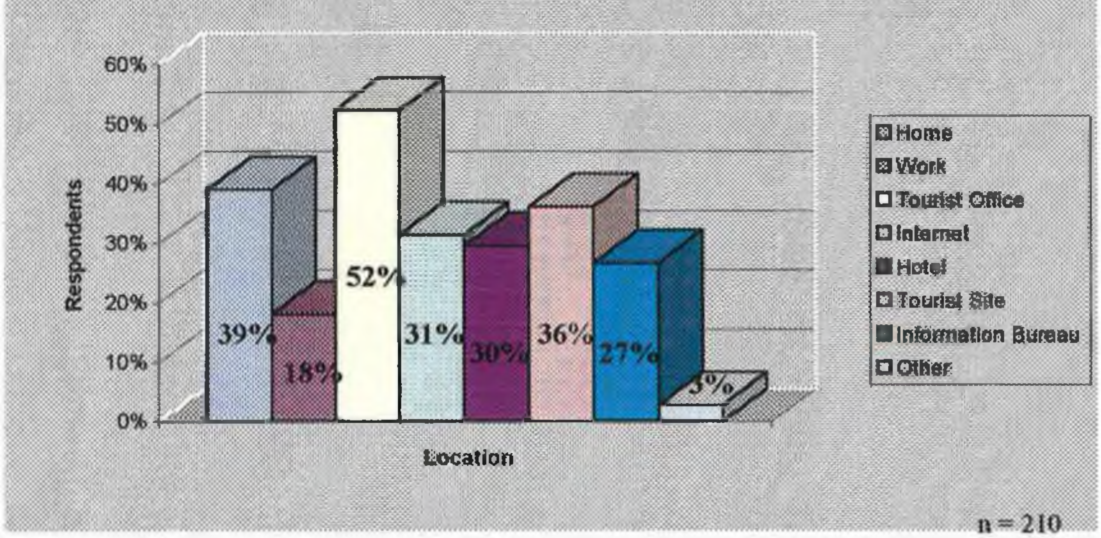

A total of 39% of the respondents would use such an application from their homes if the opportunity were made available. This piece of data was further analysis by cross-tabulating it against the result of how many people had access to a computer at home. The results of this cross-tabulation are illustrated in

Appendix M. The results were that only 68 out of the 82 respondents who would use D-VR3 at home actually had access to a computer at home. Some conclusions can be drawn from the remaining 14 respondents who did not have access to a computer in their home and would like to use D-VR3 at home. These conclusions are that they intended to get a computer, that they did not realise that a computer was necessary to use D-VR3 or that if they had a computer they would use D-VR3 at home.

A total of 18% of the respondents stated that they would use D-VR3 while at work. This result was further analysed by the use of cross-tabulation with the results obtained from Question 2, Appendix N, about who had access to a computer at home. While 38 respondents would use D-VR3 while at work only 24 of these respondents actually had access to a computer while at work. The problem, therefore, arose with those 14 respondents who would like to use D-VR3 at work yet did not have access to a computer at work at the time the questionnaire was administered.

The most common location that was chosen to use D-VR3 was in a Tourist Office where 52% of the respondents would be likely to use such an application. This was a significant discovery in that over half of the respondents were willing to use this type of application in a Tourist Office which, in many cases already have computers capable of running such an application without any extra expense. These findings were further analysed by cross tabulating them against the results obtained from Question 17 - the nationality of each respondent. From this cross-tabulation, Appendix O, it became quite evident that the use of Tourist Office was more commonplace among certain nationalities such as the Irish, French, Greeks, Australians and the Swedish.

The use of providing D-VR3 over the Internet was also examined. This was one area of the thesis that people found fascinating - the ability to walkthrough any tourism location and interact with other individuals in that location over the Internet from anywhere in the world. A total of 31% of the respondents would be likely to use D-VR3 over the Internet. These results were further analysed by cross-tabulating them against the results obtained from the gender, age group, tourist type, occupation and nationality questions and no significant trends were found.

The final four categories analysed were the likely use of D-VR3 in Hotels, at the tourist site, at information bureaux and, finally, at other locations. A total of 30% of the respondents would be likely to use D-VR3 in hotels, 36% would use D-VR3 at the actual tourist site, 27% of the respondents would use such an application at an information bureau and, finally, 3% would be likely to use D-VR3 in another location. The other locations that were specified in the questionnaires were at a bus/train station and supplied on a CD-ROM Disc. These were two excellent additions to the list of locations even though producing this type of application on a CD-ROM does not really constitute a location. The results of the above four categories were further analysed against the results obtained from Question 13 to Question 17 inclusive. Again no significant trend was noted in any of these cross tabulations.

Question 7 – Would this type of program make you more likely to, equally likely to, or less likely to visit a specific location? The aim of this question was to determine whether this type of application would make a potential tourist more likely, equally likely, or less likely to visit a location featured in D-VR3. The results of this question are clearly illustrated in Figure 6.9. A total of 2% of respondents would be less likely to visit a location after experiencing that location within D-VR3, 38% would be equally likely to visit, and 60% of the respondents would be more likely to visit the location. These are extremely important findings as a total of 98% of the respondents would be equally likely or more likely to visit a location that they have experienced using D-VR3. This is very important because this question was asked in order to gauge how many of the respondents viewed this type of application as a source of tourist information and how many viewed it a substitute for tourism. D-VR3 is a Tourist Information System developed to provide a potential tourist with a realistic impression of a specific tourist location and, thus, helping them to make an educated decision with regards to their tourism destination. The idea is to provide the tourist with a realistic impression of a location and, therefore, allow them to develop a realistic expectation of the tourism location.

Figure 6.9. Likelihood Of Visiting A Location After Using D-VR3.

These result were cross-tabulated against the gender, age groups, tourist types, occupation and nationalities of all the respondents but there were no noticeable trends discovered.

Question 8 – How does the program compare to traditional types of tourism information? This question compared D-VR3 to the traditional types of tourist information. The results, graphically illustrated in Figure 6.10., show that D-VR3 compares extremely well to the traditional types of tourist information. Only 1% of the respondents preferred traditional types of tourist information to D-VR3, 22% liked both the traditional methods of tourist information and D-VR3 equally, 25% of the respondents believed that D-VR3 was a slightly better form of tourist information and, finally, 52% of the respondents believed that D-VR3 was a much better source of tourism information than the traditional types. These results again reflected the enthusiasm that D-VR3 generated over the three-day questionnaire administration. These results were again cross-tabulated against other results obtained and no significant trend was uncovered.

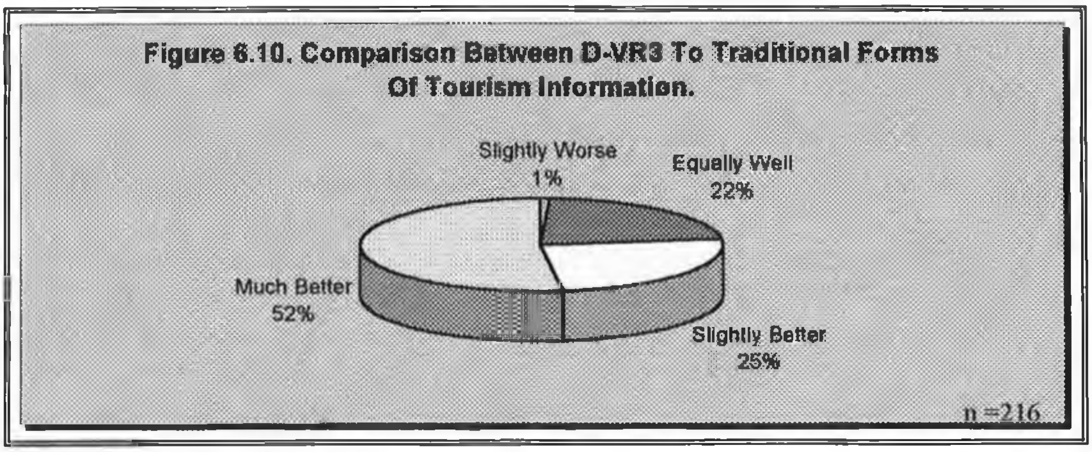

Question 9 – Have you visited the monastic enclosure in Glendalough? This relatively straightforward question was the first part of a three-part filter question. The aim of this question was to determine whether the respondent had visited the Monastic Enclosure at Glendalough. If the respondent had visited the Monastic Enclosure, then they were requested to answer the next question and if they had not, they could proceed to Question 13. The results to this question, illustrated in Figure 6.11., were that 60% of the sample had visited the Monastic Enclosure and the remaining 40% of the respondents had not visited the site yet. These results were simple and did not require any further analysis.

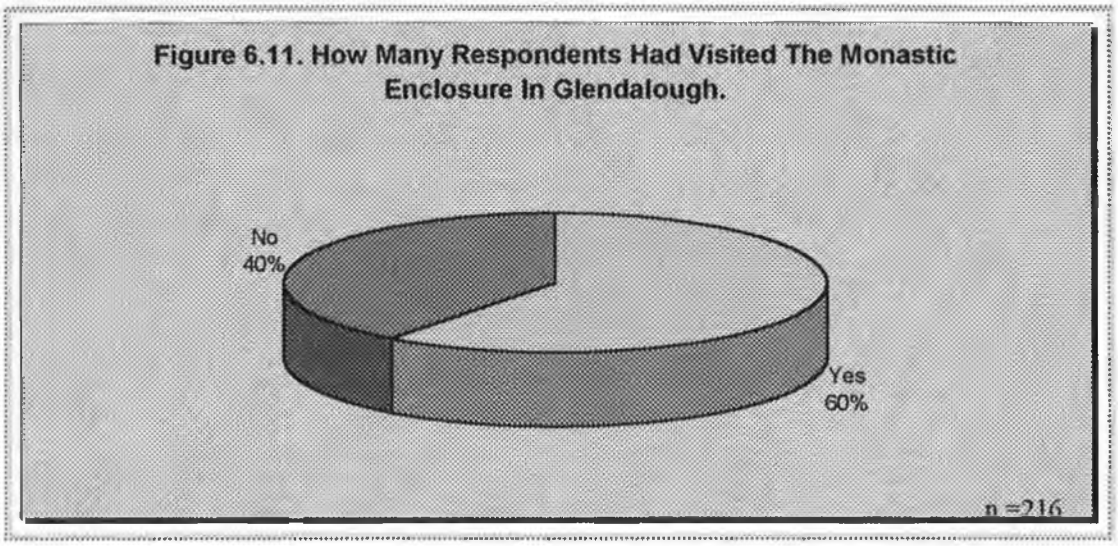

Question 10 – How do you feel that the program compares to the actual buildings? This question was designed in an attempt to determine how each respondent felt the VEs constructed within D-VR3 compared to his real-world counterparts. This was a method to gauge how each respondent felt about the realism of each VEs. Only those who responded positively to the previous question were requested to answer this question. Therefore, 40% of the sample were automatically excluded from answering this question. Of the remaining 130 respondents who did answer this question, 2% felt that the VEs in D-VR3 were much worse than the real life buildings in Glendalough, 29% believed that the VEs were only slightly worse than their real life counterparts. The remaining 69% of the respondents believed that the VEs in D-VR3 compared equally well or better to their real world counterparts. This 69% broke down into 31% of these respondents believing D-VR3 compared equally well to the real building and 20% and

18% of these respondents believing that the VEs were slightly better and much better respectively. These results are illustrated in Figure 6.12.

Figure 6.12. A Comparison Of D-VR3 To The Actual Buildings.

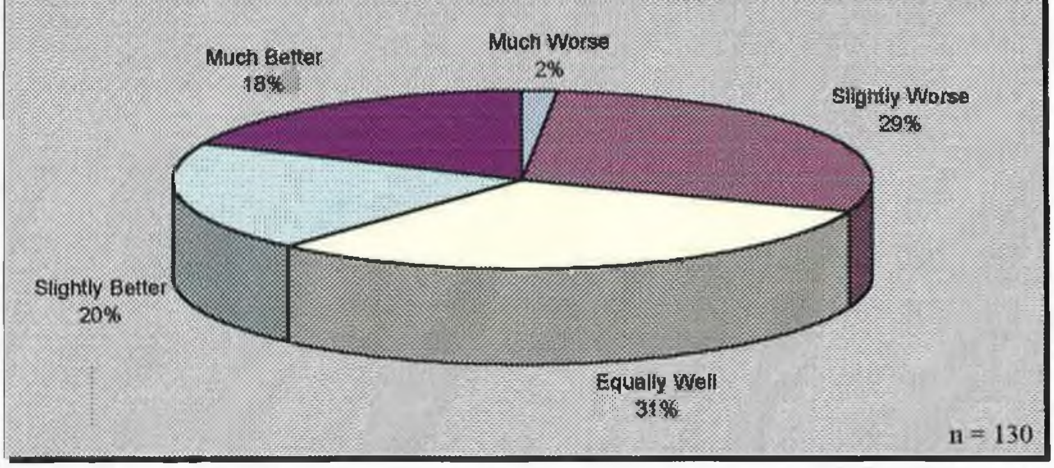

The reasons that participants felt that the VEs were slightly or, indeed, much better than the real buildings could be because the environments were reconstructions of how the buildings originally appeared and participants could view inside some of buildings which are no longer standing or no longer accessible to the general public.

Question 11 – Which model would you be more likely to use? This question was aimed at determining which of the two models, both of which were clearly indicated during the accompanying presentation, each respondent would be more likely to use. This question did not have a 100% response rate as 13% of the sample did not respond to it. Of the 188 respondents 53% would be more likely to use the first model which was a VE constructed using Virtus Walkthrough/VRML. Therefore, the remaining 47% of the respondents were more likely to use the second model, which was constructed using the 3D Realms Engine. These results, graphically displayed in Figure 6.13., were very interesting in that these two models although they were reconstructions of the same buildings were very different in the way they allowed a participant to interact with them and in the speed at which this interaction took place. The breakdown of the two models was very close as there was very little difference in these two percentages. This shows that both of these models created were very much in line with one another. This question was a filter question and only those who answered this question were requested to answer the next question. The results obtained from this question were further analysed by cross-tabulating them against each respondent's personal data obtained from Question 13 to Question 17 inclusive. No trends were uncovered from these analysis.

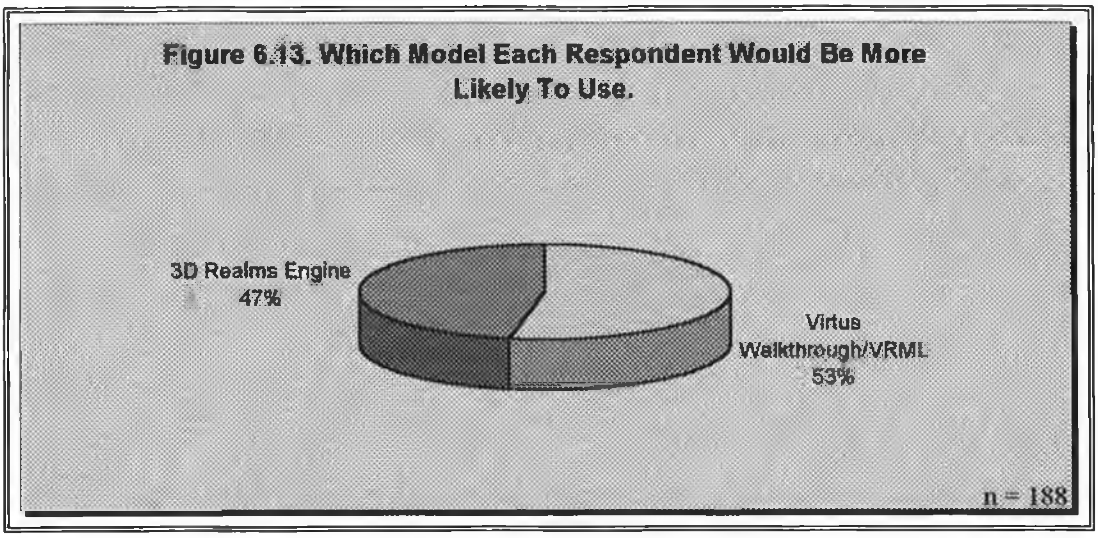

Question 12 – Why would you be more likely to use this model? This question was aimed at establishing for what reasons respondents would be more likely to use one VE instead of another and, therefore, determining what criteria a respondent would use to evaluate one world against another. An open-ended question was used so as not to influence the respondents in any way. The different criteria that respondents used in order to choose which VE they would be more likely to use are shown in Figure 6.14. Only 186 respondents completed this question from which 32% choose a specific world because they believed that it was of higher quality than the other world. A total of 23% of the respondents believed that user-friendliness was the major factor in deciding which world to use, 18% of the respondents believed that realism was the important factor in deciding which world to choose. Those respondents who believed that speed was the major factor to consider when choosing a VE constituted 13% of the respondents. Presence was believed to be the major factor in choosing which world to use by 9% of the respondents, 4% believed that more time to examine the worlds was the biggest deciding factor and, finally, 1% of the respondents believed they would choose a world depending on the time that each world take to load. The results obtained from this question were very enlightening as many of the respondents indicated what criteria they would use to distinguish which VE they preferred. All of these criteria are mirrored by the criteria used in the internal evaluation of these worlds, section 6.2. to 6.2.2. inclusive.

These results were cross-tabulated against the results obtained from Question 11 and Questions 13 to 17 inclusive. Cross tabulating these results against each respondent's personal details (Question 13 to 17) did not uncover any outstanding trends at all but when the results were cross-tabulated against the results from Question 11, Appendix P., certain significant trends appeared. A total of 54 respondents, out of the 60 who stated that higher quality was the major factor to be considered when choosing a VE, preferred worlds created using Virtus Walkthrough/VRML. Out of the 24 respondents who considered speed to be the major factor, a total of 22 preferred worlds created using the 3D Realms engine. Out of the 42 respondents who considered user friendliness to be the most important factor when choosing a VE, a total of 38 respondents preferred worlds created using the 3D Realms engine. Out of a total of 16 respondents who thought that a sense of presence was the most factor, 12 respondents preferred worlds constructed using the 3D Realms engine. Finally, out of the 34 respondents who felt that realism was the major factor to consider when choosing a VE, a total of 30 of these preferred world created using Virtus Walkthrough/VRML.

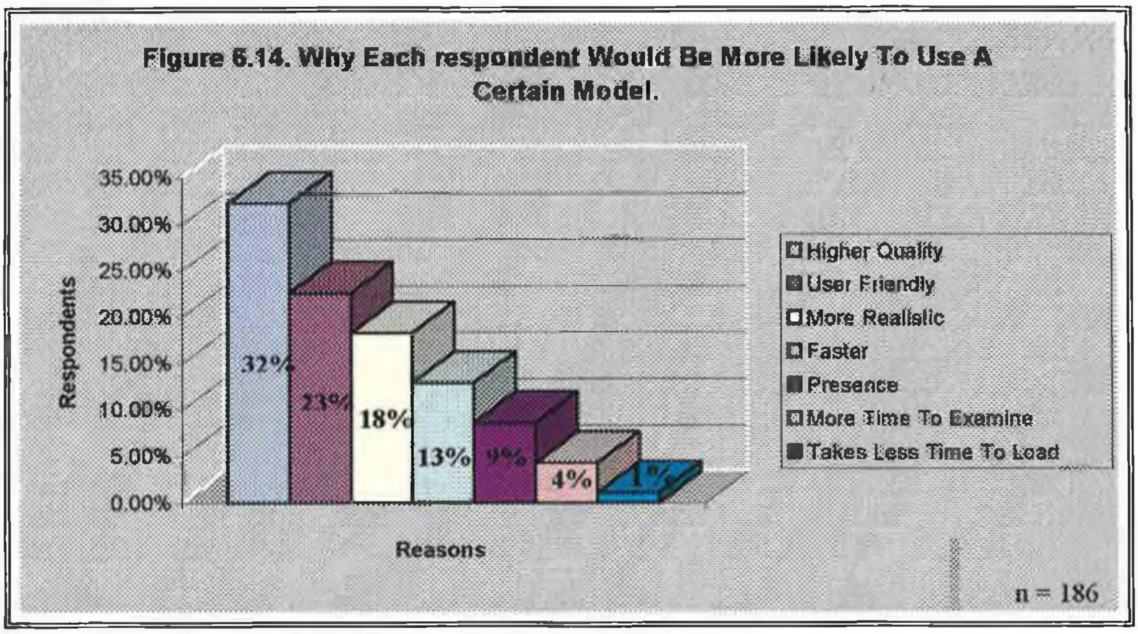

Figure 6.14. Why Each respondent Would Be More Likely To Use A Certain Model.

From this analysis it can be concluded that respondents believed that worlds created using Virtus Walkthrough/VRML were of higher quality, and more realistic than those worlds created using the 3D Realms engine. Worlds created using the 3D Realms engine, on the other hand, were considered faster and more user friendly while also providing the participant with a better sense of presence.

It is in this section where the most similarities can be drawn between the internal evaluation and the external evaluation. Both internal and external evaluation have concluded that worlds created using the 3D Realms engine are faster and provide the participant with a better sense of presence.

Question 13 to Question 17 – These questions obtained the respondent's personal details? These questions were included for two reasons. The first reason was to determine if the sample tested have the same properties as the population in general and secondly, they were used to cross-tabulate against specific questions in order to determine if any trends arose.

Question 18 – Any additional comments? This question was designed to extract any additional comments about the application or about the subject area being examined. These comments are shown in Figure 6.15. This question was only answered by 74 respondents, 34% of the total sample. The most common comments made were that the application was a very good idea, 19%, that it would be excellent to use D-VR3 as an educational tool, 16%, and that D-VR3 has great potential as a Tourist Information System, 11%. Other noticeable comments included that D-VR3 would be excellent in the case of bad weather and disabilities, 8%, it could provide a great marketing tool for most industries, 8%, and that input devices such as touch screens monitors would be very useful, 5%. All of the other comments are displayed in Figure 6.15.

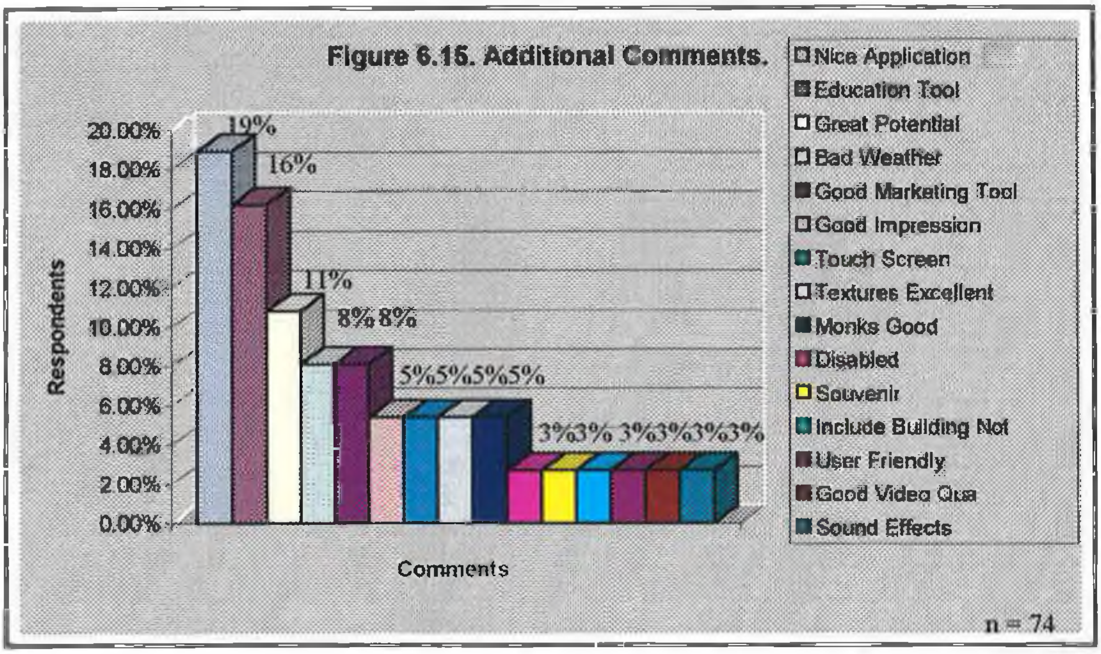

Figure 6.15. Additional Comments.

6.4. Summary.

This chapter has analysed and discussed in detail the data obtained from the internal evaluation of the VEs and the external evaluation which entailed both peer reviews and the questionnaire process used as part of the primary research. While some conclusions were made with regards to the particular findings and possible reasons for them, the remaining conclusions, along with recommendations relevant to these findings remain to be dealt with in Chapter seven.

Chapter Seven

Recommendations and Conclusions

7.1. Introduction.

This seventh and final chapter concludes this dissertation. It sums up the findings of both the primary and secondary research undertaken. Conclusions are drawn from these findings and recommendations are made concerning the role and use of VR in the tourism industry on the basis of findings uncovered in this dissertation.

7.2. Recommendations for Further Studies.

In this section several recommendations for further studies which have come to light during the course of this thesis are discussed.

7.2.1. Multiple Senses.

The Virtual Environments (VEs) designed in this thesis rely on only two senses - sight and sound. In the real world, a person's perception of his surroundings is strongly influenced by sight, sound and touch and, to a lesser degree by taste and smell. Interfaces and peripheral devices which display such stimuli are being developed and would be a welcome addition to any application as they would provide the participant with a more realistic Virtual Environment (VE). This area of research is essential for the further development of VR.

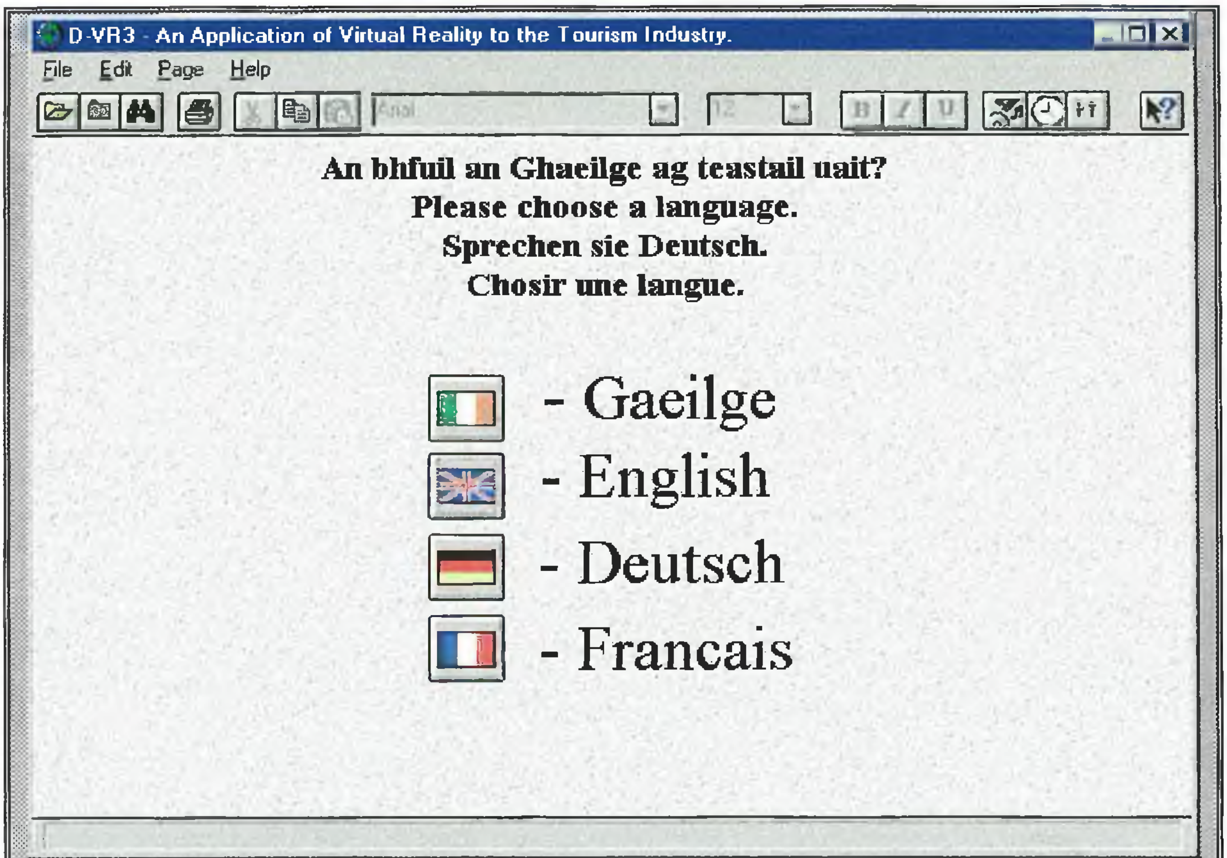

Figure 7.1. The Language Menu in D-VR3.

7.2.2. Language.

The use of language as voice communication with the VE was used in D-VR3 but only to a limited degree. The use of voice communication could be very useful as an information provider in any VR application but more so when designed from a tourism perspective. Voice communication was used in D-VR3 to provide participants with information, usually of a historical and architectural nature, but this could be widened to encompass information on open hours, prices, extra facilities and contact numbers. Care must be taken when including voice communication as one wants to convey the message clearly without causing information overload. D-VR3 also included a startup menu which allowed a participant to choose the language he wished to conduct the session in. The only language that was actually implemented was the English language. A structure was developed to enable easy input of the other languages, shown in Figure 7.1, or indeed other languages altogether. This structure also enables the appropriate sound files to be added. The inclusion of more oral and written language would enrich the experience of the participant.

7.2.3. Irrational Hyperlinks.

Hyperlink technology has been used to significant advantage in this thesis to link certain pages to other pages in the multimedia presentation and more importantly to link each VE to the multimedia presentation. However, when using certain engines, there is a one-to-one mapping of the links and the worlds they connect. This means that a link is only connected to one other link. This type of link is referred to as a rational hyperlink. Hyperlink technology does not limit a designer to rational hyperlinks. Irrational hyperlinks allow a user to enter a specific world the first time the link is encountered and to enter an entirely different world the second time that the link is encountered. These linkages could either be random or preprogrammed. Irrational hyperlinks were experimented upon in the course of this thesis but were not used in any of the final models because links used in D-VR3 were designed to be as realistic and as logical as possible. This is not to say that such technology would not be very useful in other applications of VR to the tourism industry.

7.2.4. Multiple Participant.

Although many of the VEs created during the course of this thesis have been constructed for use with multiple participants, the number of participants is limited to only a few participants, see Table 7.1. In the case of VR applications to the tourism industry the use of multiple participants is essential to allow potential Tourists to liaise with other potential tourists or with a specific person placed in the VE that would provide them with information, answer their questions or lead them on a guided tour of the facilities on offer. Therefore, the addition of multiple participants may be extremely useful as another source of information to the potential tourist. Further studies are required to increase the number of participants allowed into a specific environment at the one time. The potential of these studies to the tourism industry should also be assessed.

Table 7.1. Multi-Participant Capabilities of the Environments Evaluated.

	Modem	IPX	Serial/Parallel
Id Engine	4	2	2
Raven Engine	8	2	2
3D Realms Engine	8	2	2
Virtus Walkthrough	N/A	N/A	N/A
VRML	Depends	On	Server

7.2.5. Improvement of Input Devices.

VR equipment often does not allow free movement to the participant. A direct consequence of the restricted freedom of movements is the problem of how the participant maneuvers through a VE. In an attempt to make the application as effective and as user-friendly as possible the input devices need to be improved significantly. This issue has already been addressed during this dissertation but requires further consideration.

7.2.6. On-Line Reservations.

There is a definite trend in the tourism industry focusing on direct real-time on-line 'seamless' booking in which a potential tourist can reserve any part of the tourism product directly through an on-line booking facility. Progression of D-VR3 from being an information provider to being a complete on-line booking system is another issue which needs further consideration.

7.2.7. Expand to the Hospitality Industry.

In addition, to expanding D-VR3 to include more tourist sites, there is great potential for such an application in the Hospitality Industry. Potential tourists could not only be provided with information on the tourist attractions in a given area but would also be provided with sufficient information to chose appropriate accommodation, restaurants and facilities in that area. They could then select a source of accommodation and walk through the establishment before they choose which establishment would suit their needs. The potential tourist could then plan his trip and make the necessary reservations through an on-line reservation system such as the example described in section 7.2.6.

7.2.8. Customisation.

The application developed in this thesis has been constructed in such a way as to enable the layout to be easily configurable and customisable by the designer. This option was included in the application because further customisation of software, interfaces, and even the contents of the VEs may be required to enable the participant to configure the environment to his personal tastes. Although these capabilities were included in this application further studies need to be conducted to help to make this customisation and configuration more automated and user-friendly.

7.2.9. Database Access.

The addition of a database system to any application of VR to the tourism industry could be extremely useful for two reasons. Firstly, potential tourists could gain access to a vast amount of information about a particular destination and secondly, specific information about a potential tourist such as his identity, contact number and specific area of interest could be monitored and recorded into a database and used for marketing purposes.

7.2.10. Embodiment.

User embodiment means the provision of appropriate body images to participants so as to represent them to others and also to themselves in the VE. The use of appropriate body images in a VE provides immediate and continuous information about a participant's presence, activity, attention, availability and location. Therefore, embodiment becomes an important issue when designing multi-participant VEs. The issue of embodiment is one which is very often neglected by designers as it appears many designers still view VEs from the outside looking in rather than an environment in which one actively participates. The primary goal of a Virtual Body is to indicate a participant's presence in a VE and, therefore, where possible in the VEs in D-VR3 each participant was supplied with a human figure, Figure 5.13. However, recognising who a participant is from his embodiment is another key issue which needs further consideration with reference to D-VR3.

7.3. Conclusions.

Two words best characterise the future of the tourism industry; growth and change. The tourism industry is a major growth sector with great potential and, with an estimated 5% growth rate over the next decade, is poised to enjoy future long-term growth (WTTC 1996). International tourism has risen to 11% of the world's GNP and shows little sign of slowing down (WTTC 1996). The World Travel and Tourism Council (WTTC) have estimated the gross output of the tourism industry to be US$ 3.6 trillion in 1996 (WTTC 1996). By the year 2006 this figure is expected to have almost doubled to US$ 7.1 trillion (WTTC 1996). However, the tourism product is largely intangible, fragmented, perishable, heterogeneous and volatile and unlike other more tangible products it cannot be sampled prior to purchase. Therefore, information is often the only way that the tourism product can be presented to the potential tourist. It is essential that this information is timely, accurate and relevant as potential tourists often rely on a wealth of information before choosing a tourism destination.

Traditional sources of tourism information, images, text, sound, animation and video, only provide potential tourists with short and often limited glimpses of tourism destinations. The sources often do not provide enough information to enable a potential tourist to make an informed choice with regards to their tourist destination. In addition, these sources of tourist information only provide a passive experience as they often possess no involvement on the part of the potential tourist. This, to an extent, limits their effectiveness as a means of encouraging a potential tourist to a particular destination.

VR, on the other hand, enables a potential tourist to interact with and experience each tourist destination in high detail and provides his with enough information to make a well-informed tourist decision. When considering its application within the tourism industry, VR will offer the ability not only to view a destination but also to participate in the activities offered at the destination. Through VR the tourist advances from being a passive observer to being an active participant (Williams & Hobson 1994). VR at present can be applied to the tourism industry in two main areas - tourism policy and planning and the marketing of the tourism product. Most of the exciting developments in destination and product presentation are relevant at the information stage, within the search process for a suitable holiday destination. It is in the area of the provision of information and the marketing of the tourism destination that VR is most useful at present. Therefore, D-VR3, the application of VR designed in this thesis, is a Tourism Information System (TIS) which includes the traditional sources of tourism information but also incorporates real-time, interactive and highly vivid VEs which provides a potential tourist with the opportunity to walk through the tourism destination at their own leisure as well as interact with their surroundings. It was decided to use commercially available PC's with no additional hardware in an attempt to make the application as diffussable as possible and in order to bypass the limitations which still exist with immersion technology. One of the first conclusions to be drawn from this research is that real-time interactive and realistic walkthroughs of areas of tourist interest are possible on commercially available PCs and that the interactivity that VR provides is probably the key differentiating factor between VR and other types of tourist information. The fundamental question for tourism practitioners and researchers is whether these systems will benefit both the industry and the economy of the destination countries and help the consumer to make a more satisfactory choice (Sussmann 1994). The survey conducted during this thesis indicated that 77% of the respondents thought that VR as a source of tourist information in conjunction with the traditional sources of the tourist information was either slightly better or much better than traditional types of tourist information. It can be concluded that this 77% though that VR improved the information provided either slightly or considerably.

Applications of VR to the tourism industry will, therefore, help to generate realistic impressions and expectations of what to expect at such a destination and, thus, provide potential tourists with extra information to make an informed decision on their tourism destination. Therefore, the underlying significance of VR is its ability to bring the experience to the customer and this virtual experience of the destination will further increase the customer's desire to actually visit the location (Coates 1992). There are, however, many who believe that this virtual experience will ultimately act as a substitute to the tourism product (Cheong 1995). As convincing as the arguments supporting the eventual replacement of tourism by VR might seem, there are many (Musil & Pigel 1994) who adamantly insist that VR can never serve as a substitute to tourism. Musil and Pigel maintain that VR can never replace the feeling of being in nature and seeing, hearing, feeling and breathing an environment that is real. Furthermore, the real world possesses a high level of complexity, randomness and uncertainty which VR systems might not have the ability or speed to handle. Furthermore, these systems may not be able to process the vast amount of information required to project a realistic virtual experience (McClure 1994). In the case of the primary research undertaken potential tourists were either as likely, or more likely, to visit a location after experiencing a VE. A total of 60% of the respondents would be more likely to visit a location and 38% would be as likely to visit a specific location after using D-VR3 or a similar application. Therefore, it can be concluded from this research that a vast majority of the respondents, 98%, viewed this application as a decision making tool which increases their desire to actually visit a tourist location and, therefore, did not view the application as a tourism substitute.

VR is a significantly compelling creative environment as a source of tourism information but certain improvements in technology are necessary in order improve the realism of the VEs. In general, respondents were fascinated by the experience of the Virtual Worlds but improvement in sensory stimuli is essential if VR is going to advance but interaction in VR is also vital to give a sense of presence (Rheingold 1991). The first two sensory stimulations, sight and sound, used in D-VR3 and the third sensory input, touch, used in some Immersive VEs, are at present rather crude and unrefined virtual stimulations. The image quality, resolution and movement have yet to achieve adequate levels of realism owing to limitations in processor power and speed. The remaining two stimulations, smell and taste, have yet to be mastered. The development of faster and more powerful processors in the future will overcome some of these shortcomings and allow more realistic representations of the Virtual World. The quality of visual and auditory stimulation may eventually be improved to levels indistinguishable from those of our world, while tactile stimulation will be conveyed to all parts of the body via body suits (Cheong 1995). However, from a consumer's perspective, group participation, control and intellectual stimulation will be more important than sensory overload. These technologies will create even more realistic VEs which in turn will be more likely to speed up VR's commercial availability (Maloney 1992) and its acceptance.

In all, several uses and potential uses of the application in the tourism industry were uncovered. Most of these applications were based on the belief that D-VR3 provided the potential tourist with a good perception of what he would experience if he were to visit a specific tourist location. Not only does D-VR3 provide this perception and some general information about the site but it also provides the user with ancillary historical and architectural information. The combination of the perception of what is available at the site and the different types of information provided by D-VR3 make it an extremely useful tool for deciding on which site to visit and planing one's visit once one arrives at the site. The places where respondents would be most likely to use an application of VR such as D-VR3 would be in a tourist Office, in their own home or at the actual tourist Site. D-VR3, or a similar application, would be supplied on a CD-ROM or the Internet so that respondents could use the application at the location they find most convenient. Although until recently consumers were reluctant to utilise home holiday purchasing (Gilbert 1990) the increasing computer literacy of younger generations, the wider availability of IT in more households and the growth of IT's capabilities are expected to increase the direct involvement of consumers with IT, in order to get a wider range of information on destinations and their attractions before making their choice (Buhalis 1994).

With regard to the software used to create VEs, results from the external evaluation show that 53% of the respondents would be more likely to use VEs created using Virtus Walkthrough/VRML compared to 47% of the population which would use environments created using the 3D Realms engine. These results indicated that neither engine was viewed as being far superior to the other and when these results were viewed in conjunction with other finding the respondents who chose VEs created using Virtus Walkthrough/VRML did so because they perceived these worlds as being of higher quality and detail than those created using the 3D Realms engine. Respondents who would prefer to use VEs created using the 3D Realms engine did so because they preferred faster and more user-friendly environments which provide them with a better sense of presence. The results from the Internal evaluation further reinforced the results from the external evaluation. The internal evaluation measured each VE against a strict set of evaluation criteria which were grouped into three main sections: interactivity, vividness, and time. No one criterion can be examined in isolation but instead all criteria must be examined in conjunction with all the other criteria. It can therefore be concluded that although real-time, interactive and highly vivid VEs of areas of tourism interest are possible on commercially available PCs it is still too soon to predict if one particular piece of software is the most appropriate for all tourism sites.

To conclude, although VR is still a relatively novel concept, VR has already demonstrated its capabilities and potential in many industrial and scientific fields. The reality is that VR technology has a long way to go. It has already been noted that present sensory stimulation has yet to attain realistic levels but it is hoped that future advances in computer technology will remedy these shortcomings. In addition, problems which plague other computer-based systems will continue to persist. Before the potential of VR can be fully realised, such limitations will need to be fully addressed.

The unique features of VR provide a revolutionary way to visualise and interact with tourism information, but at present it is impossible to ascertain or extrapolate with any degree of confidence the possible impact that VR will have on the tourism industry. The natural pressures and tendencies to oversell, by both the media as well as the practitioners, may threaten a backlash in public enthusiasm and financial support from funding sources. When dealing with the development of VR it is still in its infancy and certain promises may not be fulfilled for quite some time.

However, as technology costs decrease and computer power increases, the development of VR experiences is coming much closer. The quality of the computer-generated visual images has improved enormously over the last few years, with graphics, animations and digitised video images having clearer resolution than ever before. This imaging technology is continuously improving and will play a key role in future VR development. As the quality of visual images and sound improves, the degree of realism in Virtual experiences will be enhanced. Its intensity, along with its ease of use, will make VR a potent technology. In its own way it has the capacity to effect society as deeply as the telephone or television. Both have changed the way we communicate with our family and friends, changed how we view the world around us, and changed the way we approach business, politics, religion and personal relationships. From home entertainment to tourism and from flight simulators to surgical simulators, VR will play an increasingly important role in public and private life as we move towards the twenty-first century.

Glossary

3-D	three dimensional; a display, medium or performance giving the appearance of height, width and depth.
6-DOF	six degrees of freedom; six different measurements that can be assigned to any one movement. E.g., any one movement may entail changes in as many as three spatial positions (horizontal, vertical and depth) and three degrees of orientation (rotation, yaw, and pitch). The capability for assessing all six for each movement is incorporated into many devices, such as six-degree-of-freedom tracking mechanisms.
AI	Artificial Intelligence; computer programs that try to efficiently perform task-orientated computations (e.g., memorising lists) or to stimulate cognitive aspects of human behaviour (e.g., figuring out which lists to memorise); an attempt to represent, process and transfer knowledge.
Artificial Reality	simulated, computer-generated spaces; a combination of computer and video systems; term coined by Myron Krueger for title of groundbreaking book.
Author (VR)	designer of programs of Virtual Environments. An expert enlisted from an outside discipline is sometimes called a collaborative author.
Bandwidth	the particular range of electronical frequencies (consequently, the types of communications) that can be handled by a communication device.
CD-ROM	Compact-Disk Read-Only Memory storage device; can hold a book-full of text ; valuable delivery media but store no information about/by users. Writable optical type is called WORM (Write-Once Read Many times).
CPU	Central Processing Unit; the main part of the computer that interprets and execute commands as they are received; part containing logic and computational circuitry - sometimes all on the one chip.
CRTs	Cathode-Ray Tubes; around since the end of the 19th Century; glass display terminal or TV screens, heavy and bulky for many purposes, sometimes replaced by LCDs. Images are produced by electrons shot into a light emitting phosphor screen.
Cyberspace	term coined by William Gibson, science fiction writer to describe computer-synthesised, 3-D spaces; fantasy site of "consensual hallucination." The term "Cyberspace" and "VR" are often interchangeable, but some researchers make a distinction between them based on usage or equipment requirement. Most often VR involves "experiences," and Cyberspace is used more for visualising information

and accessing it.

Cyborg	robotic humanoid modelled directly from digital reading of a real actor and transformed into a photorealistic, animated "twin" such as those shown in the action films *Terminator 2* or *Robocop*; spiritless character produced via illusionary metamorphosis.
Environment	space in which a user of VR technology imagines him- or herself and within which interactions take place; computer-generated display world or scenario.
Expert System	computer program relying on knowledge or reasoning to emulate the performance of a human expert; contains encoded rule statements (in a knowledge base or rule base) that reflect individual or gather expertise in a field and then performs "reasoning" via a rule interpreter in order to reach a decision, come to a conclusion or give up (if unable to do either of the former).
Fiber-Optic	sending large amounts of data as light pulses through slender cable.
Haptic	pertaining to sensations of touch, pressure, temperature, twist, etc., mediated by skin, muscle, tendons, or joints, representing tactile senses (though not limited to these).
Haptic Displays	computer-generated display designed to appeal to the haptic senses i.e., one in which things seem to push or pull, radiate or sensate in different directions with varying degrees of strength; objects within the virtual world with assigned force fields, torque, friction, heat and pressure, which are made noticeable to a person into acting with them.
HDTV	High Definition Television; system for transmission of a very high resolution television signal.
HMD	Head Mounted Display; headpiece or head-held brace with viewing or optical devices (located or suspended in front of the user's eyes).
Hypermedia	information in the form of text, diagrams, animation, images, moving images, speech, sound or computer programs; information under the control of a computer and the means for a user to navigate through it in valuable ways. Hypermedia are frequently used for complex documents , such as training manuals, safety manuals and other large documents stored on optical disks.
I/O device	Input/Output device used to present or receive data via computer.

Immersive Graphics encompassing multidimensional displays (pictures, sound, tactual effects) wrapped around so that all senses seem immersed and the line between the real and the illusionary worlds disappears.

Interactive — has features enabling a user to influence or manipulate the course of action; allows conversant interdependency between user and system.

Interface — whatever is used - a piece of equipment or choice of menus to connect a user with a computer program; whatever is used (often an interface card inserted into the computer unit) to connect a computer to an external such as a printer, modem or network.

LEDs — Light-Emitting Diodes; in VR applications, little things used as beacons in combination with imaging sensors (usually mounted on headgear) for tracking purposes; also placed at the end of joint sections in instrumented gloves to aid measurement of the degree of bending.

MIDI — Musical Instrument Digital Interface; a standard code that accepts as commands different types of signals, including bio-controller signals, which it then implements to control electronic instruments, e.g., music, synthesisers, word processors, household devices, etc.

Model (VR) — graphics databases of the colours and coordinates for polygon-shaped pieces of the virtual world; a computer-generated simulation of something real.

Multimedia — combined text, images, full-motion video and sound; requires lots of bandwidth and computing power.

NASA — National Aeronautics and Space Administration; a division of the U.S. federal government.

Paradigm — an example or conceptual model used to illustrate new approaches or ways of thinking.

Parameter — measurement factors or determined bounds.

Pixels — short for picture elements; dots on a computer screen that makes up letters or pictures. The number of pixels per inch determines the sharpness of the image.

Power Glove — VPL version of DataGlove for home use.

RB2	Reality Built for 2; Reality Built for 1 (RB1), VPL product of the 1980's, one of the first commercial, real-time, on-line, single user systems; allows users to interact with virtual objects. Reality Built for 2 follows the same concept as RB1, but is for two people. It allows two users to interact not only with objects but also with each other.
Real-time	the actual time something occurs; for problem-solving with computers, the time between putting data in and receiving a solution; used when response to input to a computer is fast enough to affect subsequent input.
Rendering	translation into another form, e.g., converting signals into a picture; yielding or reducing to another state or interpretation.
RGB	the primary colours Red, Green, Blue; additive colours used by colour monitor displays. The combination and intensities of these three colours are used to represent the whole spectrum.
Speech Recognition	Recognition of the human voice as input to a computer; transcribing whatever is spoken, analysing the sound patterns and then converting them into digitised text.
Sutherland, Ivan	started the whole field of VR in 1965 with visions of "the screen as a window into which one looks into a virtual world," one that would "look real, sound real, feel real" (Sutherland, 1965); built the first head-mounted gear in Utah in 1968.
Tactile Feedback	feedback directed through or simulating the sense of touch or physical feel; sense of contact distinguished from force feedback, which reflects magnitude of force.
Teleoperation	carrying out specific tasks via robots or telepresence; also refereed to as telemanipulation.
Telepresence	word coined by Marvin Minsky; "remote" presence; medium that gives a person the sense of being placed physically within remote, computer-created scenes; a psychological experience resulting when simulator technology works well enough to convince users they are immersed in virtual worlds (Rheingold, 1991).
Tracing	the simulation of light effects.
Transmission	filaments; expensive but accurate, reliable and fast, especially over long distances.
Visualisation	taking data, exploring its meaning and making it more comprehensible by presenting it in an intuitive simulation; used primarily for physics, chemistry and medical applications.

VPL	Virtual Programming Language (Laboratory, Inc.), one of the first commercial ventures for developing and producing control devices (e.g., DataGloves) for Virtual Environments, founded by Jaron Lanier.
VR	Virtual Reality is a human-computer interface where the computer and its devices create a sensory environment that is dynamically controlled by the action of the individual so that the Virtual Environment appears real to the participant.
WIMP	acronym for widget, icon, mouse and pull-down-menu approach to graphical user interfaces.
Workstation	a single user microcomputer, usually with high resolution graphics capability and high-speed capacity, capable of running applications independently or with other computers via a network. Workstations are generally considered more powerful than personal computers (PCs); however, many high-end PCs now match or exceed low-end workstations and are less expensive. Today's PC approaches in power the large IBM mainframes of the 1980s.

Bibliography

Andreassis, **"Dentists Offer Video Diversion for Patients,"** Hampton-Richland Eagle, November 4th 1994.

Archdale, G., **"Computer Reservation Systems and Public Tourist Offices."** Tourism Management, 14 (1), 1993.

Balaguer, F., and A. Mangili, **"Virtual Environments."** In Thalmann, N.M., and D. Thalmann (Eds.), New Trends in Animation and Visualisation, Wiley and Sons, New York, 1991.

Bennett, M and M. Radburn, **"Information Technology in Tourism; The Impact on the Industry and Supply of Holidays."** In Sinclair, M.T and M.J. Stabler, (Eds.), The Tourism Industry; An International Analysis, Oxford: CAB International. 1991.

Bennett, R.P., **"Walking Virtually: Help for Parkinson's Disease."** Virtual Reality Special Report, Spring, 1995.

Biocca, F., **"The Problem of Simulation Sickness and the Diffusion of Virtual Reality Technology."** In Proceedings of the Meeting of the International Communication Association, Miami, 1992.

Bricken, M., **"Virtual Worlds: No Interface to Design."** In M. Benedikt (Ed.), Cyberspace: First Steps, MIT Press, Cambridge, MA., 1991.

Bricken, M.S., **"A Description of the Virtual Reality Learning Environment."** Human Interface Technology Laboratory Technical Report, Seattle, WA., 1990.

Bricken, M.S., **"Virtual Reality Learning Environments: Potentials and Challenges."** Human Interface Technology Laboratory Technical Report, Seattle, WA., 1991.

Bricken, M.S., and C.M. Byrne, **"Students in Virtual Reality: A Pilot Study."** In Alen Wexelblat (Ed.) Virtual Reality: Applications and Explorations, Academic Press, San Diego, CA., 1993.

Bricken, W., **"Learning in Virtual Reality."** Human Interface Technology Laboratory Technical Report, Seattle, WA., 1990.

Bricken, W., **"VEOS Preliminary Functional Design."** Human Interface Technology Laboratory Technical Report, Seattle, WA., 1991

Brooks, F., **"Grasping Reality Through Illusion: Interactive Graphics Serving Science."** Technical Report No. 88-007, Department of Computer Science, University of North Carolina, Chapel Hill, 1988.

Brooks, F.P., et al., **"Project GROPE - Haptic Displays for Scientific Visualisation."** Computer Graphics: Proceedings of SIGGRAPH, Dallas, TX., 1990.

Buhalis, D., **"Information and Telecommunication Technologies as a Strategic Tool for Small and Medium Tourism Enterprises in the Contemporary Business Environment."** In Seaton A.V., et al. (Eds.), Tourism: The State of the Art, Chichester, London, 1994.

Cheong, R., **"The Virtual Threat to Travel and Tourism."** Tourism Management, Vol. 16 (6), 1995, pp. 417-422.

Cleverdon, R., **"Global Tourism Trends: Influences and Determinants."** In World Travel and Tourism Review, Volume 2, Ritchie, J.R., and D. Hawkins (Eds.), CAB International, Wallingford, Oxford: 1992.

Coates, J., **"The Future of Tourism: The Effects of Science and Technology."** Vital Speeches of the Day, October, 1992.

Collier, D., **"Expansion and Development of CRS."** Tourism Management, 10 (2), 1989, pp. 86-88.

Coull, T.B., **"Texture-Based Virtual Reality on a Desktop Computer Using WorldToolkit."** In Proceedings of Virtual Reality, San Francisco, 1991.

Dennett, D., **"Consciousness Explained."** Little, Brown, Boston, 1991.

Digital Image FX., **"Virtual Reality Feasibility Study."** Digital Image FX, Dalhousie, 1995.

Earnshaw, R.A., et al., **"Virtual Reality Systems."** Academic Press Ltd., London, 1995.

Ferrington, G., and K. Loge, **"Virtual Reality: A New Learning Environment."** The Computing Teacher, 19 (7), 1992.

Furness, T.A., **"Visually Coupled Information Systems."** ARPA Conference on Biocybernetic Applications for Military Systems, Chicago, IL., 1978.

Furness, T.A., **"Creating Better Virtual Worlds."** In Proceedings of Human-Machine Interfaces for Teleoperators and Virtual Environments, Washington, DC., 1990.

Furness, T.A., **"Harnessing Virtual Space."** Society for Information Display International Symposium, Playa del Rey, CA., 1988.

Gamble, P.R., **"Connectivity and Pan-European Marketing."** International Journal of Contemporary International Management 3 (4), 1991, pp. 37-41.

Gamble, P.R., **"Small Computers and Hospitality Management."** Hutchinson, London, 1984.

Gibson, J., **"The Senses Considered as a Perceptual System."** Houghton-Mifflin, Boston, 1966.

Gibson, J.J., **"The Ecological Approach to Visual Perception."** Houghton-Mifflin, Boston, 1979.

Gilbert, D., **"European Tourism Product Purchase Methods and Systems."** The Service Industry Journal, 10 (4), 1990.

Go, F., **"The Role of Computerised Reservation Systems in the Hospitality Industry."** Tourism Management, 13 (1), 1992.

Goodall, B., **"How Tourists Choose their Holidays: An Analytical Framework."** In Goodall, B., and G. Ashworth (Eds.), Marketing in the Tourism Industry - The Promotion of Destination Regions, Routledge, London, 1990.

Greenleaf, W., **"Virtual Reality: Solutions for Rehabilitation, Ergonomics, and Disability."** Virtual Reality Special Report, January/February, 1996.

Hanefors, M., and L. Larson, **"Video Strategies used by Tour Operators."** Tourism Management, 14 (1), 1993.

Hitchins, F., **"The Influence of Technology on U.K. Travel Agents."** Travel and Tourism Analyst, (3), 1991, pp. 88 -105.

Hobson, J., and A. Williams, **"Virtual Reality: A New Horizon for the Tourism Industry."** Journal of Vacation Marketing, Vol. 1, (2), 1995, pp. 125-135.

Hobson, J., and M. Usyal, **"Infrastructure: The Silent Crisis Facing the Future of the Tourism Industry."** Hospitality Research Journal, 17 (1), 1993.

Jacobson, L., **"Garage Virtual Reality."** Sams, Indianapolis, IN., 1994.

Kalawsky, R., **"The Science of Virtual Reality and Virtual Environments: A Technical Scientific and Engineering Reference on Virtual Environments."** Addison-Wesley, reading, Ma., 1993.

Kasavana, M., and J. Cahill, **"Managing Computer Systems in the Hospitality Industry."** CBI, New York, 1987.

Kotler, P., **"Marketing Management: Analysis, Planning, Implementation and Control."** McGraw-Hill, London, 1991.

Larijani, C., **"The Virtual Reality Primer."** McGraw-Hill, New York, 1994.

Latta, J., **"When will Virtual Reality Meet the Market Place."** In Proceedings at the Virtual Reality Conference, San Francisco, 1991.

Laurel, B., **"Computers as Theater."** Addison-Wesley, Menlo Park, CA., 1991.

Lovine, J., **"Step into Virtual Reality."** Windcrest/McGraw-Hill, New York, 1995..

Maloney, J., **"Are We Having Fun Yet?"** Digital Media: A Seybold Report, 2 (2), 1992.

McClure, M., **"The Travel Experience: Technology as a Threat of Opportunity?"** Information Technology for Travel and Tourism Marketing: A Tool for Profit PATA Conference, February, 1994.

Merril, J., et al., **"Surgical Simulation Using Virtual Reality Technology: Design, Implementation, and Implications."** In Szabo, Z., (Ed.), Surgical Technology International III, 1995.

Mestel, R., **"Virtual Actors Help the Medicine Go Down,"** New Scientist, September 9th, 1993.

Mill, R.C. and A.M. Morrison, **"The Tourism System: An Introductory Text."** Hertfordshire: Prentice-Hall, 1985.

Musil, S., and G. Pigel, **"Can Tourism be Replaced by Virtual Reality Technology."** In Proceedings of the International Conference in Innsbruck, Austria, Springer Verlag, 1994.

Naimark, M., **"Realness and Interactivity."** In B. Laurel (Eds.), The Art of Human-Computer Interface Design, Addison-Wesley, Reading, MA., 1990.

Norman, D., **"The Design of Everyday Things."** Doubleday, New York, 1988.

Office of Public Works, **Visitors Figures**, 1996.

Oppenheim, A.N., **"Questionnaire Design: Interviewing and Attitude Measurements."** Pinter Publishing, London, 1992.

Papert, S., **"Mindstorm."** Basic Books, New York, 1980.

Papert, S., **"Situating Constructionism."** In I. Harel, and S. Papert (Eds.), Constructionism. Ablex Publishing Corporation. Norwood, NJ., 1991.

Pimentel, K., and K. Teixeira., **"Virtual Reality - Through the New Looking Glass."** Windcrest/McGraw-Hill, New York, 1992.

Quarendon, P., **"Creating Three-Dimensional Models from Sequences of Images."** In Earnshaw, R.A., et al., **"Virtual Reality Systems."** Academic Press Ltd., London, 1995.

Rheingold, H., **"Virtual Reality."** Summit Books, New York, 1991.

Riley, R., and C.S. Van Doren, **"Movies as Tourism Promotion - A "Pull" Factor in a "Push" Location."** Tourism Management, 13 (3), 1992.

Rogers, E., **"The Diffusions of Innovations."** The Free Press, New York, 1983.

Rothbaum B.O., et al., **"Effectiveness of Computer Generated Graded Exposure in the Treatment of Acrophobia."** American Journal of Psychiatry, 152, (4), 1995.

Schmid, B., **"Electronic Markets in Tourism."** In Proceedings of the International Conference in Innsbruck, Austria, Springer Verlag, 1994.

Schmitt, G.N., **"Virtual Reality in Architecture."** In Thalmann, N.M., and D. Thalmann (Eds.), Virtual Worlds and Multimedia, Wiley and Sons, New York, 1993.

Schratt, R.D., et al., **"Virtual Reality for Improving Control in Endoscopic Surgery."** In Medicine Meets Virtual Reality II, San Diego, 1994.

Scott, A., **"Virtus Walkthrough Pro: Professional 3D Visualisation – User's Guide."** Heath Scofield, Cary, NC., 1994.

Sheridan, T., **"Musing on Telepresence and Virtual Presence."** Presence, 1 (1), 1992.

Shneiderman, B., **"Designing the User Interface: Strategies for Effective Human-Computer Interaction."** Addison-Wesley, Reading, MA., 1992.

Stenvert, R., **"Constructing the Past: Computer-Assisted Architectural-Historical Research."** Ph.D. Dissertation, Rijksuniversiteit te Utrecht, The Netherlands, 1992.

Steuer, J., **"Defining Virtual Reality: Dimensions Determining Telepresence."** Journal of Communication, 42 (4), 1992, pp. 73-93.

Stoppi, J., **"Virtual & Real-Time Interactive Spatial Modelling."** In Proceedings of the Conference on Virtual Reality International: Impacts and Implications, London, 1992.

Sturman, D.J., et al., **"Hands-on Interaction with Virtual Environments."** In Proceedings ACM SIGGRAPH Symposium on User Interface Software and Technology, Williamsburg, VA., 1989, pp. 19-24.

Sussmann, S., **"The Impact of New Technological Developments on Destination Management Systems."** In Cooper, C., and A. Lookwood (Eds.), Progress in Tourism, Recreation and Hospitality Management, Vol. 5, Wiley and Sons, Toronto, 1994.

Sutherland, I., **"The Ultimate Display."** In Proceedings of the IFIP Congress, Vol. 2., 1965.

Swarden., C.G., **"What's Showing at the Dentist's,"** The New York Times, December 11th 1994.

Vaughan, T., **"Multimedia: Making It Work."** McGraw-Hill, New York, 1994.

Vygotsky, L., **"Mind in Society: The Development of Higher Psychological Processes."** Harvard University Press, Cambridge, MA., 1978.

Wang, J., et al., **"A Real-Time Optical 3D Tracker for Head-Mounted Display Systems."** Computer Graphics: 1990 Symposium on Interactive 3D Graphics, Vol.24, No.2, 1990.

Williams, A.P., **"Information Technology and Tourism: A Dependent Factor for Future Survival."** In World Travel and Tourism Review, Vol. 3, Ritchie, J.R., and D. Hawkins (Eds.), CAB International, Wallingford, Oxford, 1993.

Williams, A.P., and J.S.P. Hobson, **"Tourism: The next generation. VR and Surrogate Travel, is it the Future of the Tourism Industry?"** In Seaton A.V., et al. (Eds.), Tourism: The State of the Art, Chichester, London, 1994.

Winn, W.D., **"A Conceptual Basis for Educational Applications of Virtual Reality."** Human Interface Technology Laboratory Technical Report, Seattle, WA., 1993.

Winn, W.D., and W. Bricken, **"Designing Virtual Worlds for Use in Mathematics Education: The Example of Experiential Algebra."** Educational Technology, 32 (12), 1992, pp. 12-19.

Wodaski, R., **"Virtual Reality Madness! 1996."** Sams, Indianapolis, IN., 1996.

WTO., **"International Tourism Overview."** World Tourism Organisation, 1993.

WTTC., **"1996/7 WTTC Travel and Tourism Report."** World Travel and Tourism Council, 1996.

Appendix A

Position Tracking Systems

Electro-magnetic trackers use sets of small induction coils in body mounted sensors that are pulsed to produce magnetic fields. The participant's position and orientation can then be tracked by reading alternating magnetic fields. The magnetic sensors determine the strength and angles of the fields. These systems represent the most practical approach to date because they offer reasonable sample rate, in a small package, at a modest cost. Currently, someone using an electromagnetic tracking device is constricted to a spherical workspace of only four to five feet in diameter. Ceiling trackers, using light emitting diodes (LEDs) extend the workspace considerably - but even these devices do not allow the workspace to exceed room size[1]. Probably the biggest problem that faces electro-magnetic position tracking is time-lag, or the time required to make the measurements and pre-process them before inputting them into the simulation engine. Users of electromagnetic trackers may also encounter noise interference or performance degradation from metallic distortion in certain environments.

Despite the problems associated with electro-magnetic trackers they are the most commonly used position tracking device. The Inside-Trak from Polhemus Incorporated., Figure A.1.A., is an example of such a tracker that can be easily integrated into most PC based Virtual Environments. The Polhemus Inside -Trak generates a small electro-magnetic field and tracks the movement of receptors within it. By placing the receptors on a participant's head and hand, the rotation and translation of those parts of their physical bodies are mapped onto their virtual bodies. The participant needs to stay within the generated field and are attached to the tracking devices with wires. The accuracy of the Inside-Trak is affected by metal and most environments have to be extensively mapped. Another magnetic system, the "Flock of Birds" system, Figure A.1.B., uses a DC magnetic field and is gaining popularity because of its lower cost. Hand movement provides signals which the computer needs to allow the operator to manipulate objects.

Figure A.1.A. Polhemus Tracking System

Figure A.1.B. The "Flock of Birds" Tracking System.

Optical position tracking. By using a video camera and a set of integrated LEDs, it is possible to track movement very quickly and with great precision. The LEDs are pulsed in sequence and the cameras image is processed to detect the flashes. Optical tracking is usually possible with minimal hardware, but the

[1] Holloway, R., et al., **" Virtual-Worlds Research at the University of North Carolina at Chapel Hill"**, In Proceedings, Computer Graphics, London, October, 1991.

management of the information gathering elements requires great attention to detail. Exact positioning of the sensors and cameras, and precise calibration of the photo-diodes to gather accurate information are crucial to the success of an application. This type of system can also be produced relatively inexpensively in large quantity.

Ultrasonic, gyroscopic and mechanical devices are also used for tracking purposes but the comprise only a very small portion of the tracking techniques used for VR applications. Ultrasonic sensors are relatively inexpensive and far faster than electro-magnetic sensors but they are sensitive to outside noises and require obstruction free paths between signaller and microphone. Gyroscopic devices are accurate and compact but are expensive and need frequent recalibration. Mechanical devices are simple and inexpensive but are physically limiting.

Appendix B

Geometric Modelling Methods

B.1. Three-Dimensional PolyLines and PolyPoints.

The simplest 3D objects are known as PolyPoints and PolyLines. A PolyPoint is simply a collection of points in space. A Polyline is a set of vectors that form a continuous line.

B.2. Polygons

The most common form of objects used in VR systems are based on flat polygons. A polygon is a planar, closed multi-sided figure. They maybe convex or concave, but some systems require convex polygons. The use of polygons often gives objects a faceted look. This can be offset by more advanced rendering techniques such as the use of smooth shading and texture mapping. Some systems use simple triangles or quadrilaterals instead of more general polygons. This can simplify the rendering process, as all surfaces have a known shape. However, it can also increase the number of surfaces that need to be rendered.

Polygon Mesh Format, also referred to as Vertex Join Set, is a useful form of polygonal object. For each object in a Mesh, there is a common pool of points that are referenced by the polygons for that object. Transforming these shared points reduces the calculations needed to render the object. A point at the edge of a cube is only processed once, rather than once for each of the three edges that reference it.

The geometry format can support pre-computed polygon and vertex normals. Both polygons and vertices should be allowed a colour attribute. Different renderers may use or ignore these and possibly more advanced surface characteristics. Pre-computed polygon normals are very helpful for backface polygon removal. Vertices may also have texture coordinates assigned to support texture or other image mapping techniques.

B.3. Primitives.

Some systems provide only primitive objects, such as cubes, cones, and spheres. Sometimes, these objects can be slightly deformed by the modelling package to provide more interesting objects.

B.4. Solid Modelling and Boolean Operations.

Solid Modeling, sometimes referred to as Computer Solid Geometry (CSG) is one form of geometric modelling that uses primitive objects. It extends the concept by allowing various addition, subtraction, boolean and other operations between these primitives. This can be very useful in modelling objects when one is concerned with physical calculations, such as center of mass, etc. However, this method does incur some significant calculations and is not very useful for VR applications. It is, however, possible to convert a CSG model into polygons. Various complexity polygonal models could be made from a single high resolution “meta-object” of a CSG type.

B.5. Curves & Patches.

Another advanced form of geometric modelling is the use of curves and curved surfaces, referred to as patches. These can be very effective in representing complex shapes. However, there is significant calculation involved in determining the surface location at each pixel, thus curve based modelling is not used directly in VR systems. It is possible, however, to design an object using curves and then compute a polygonal representation of those curved patches. Various complexity polygonal models could be made from a single high resolution 'meta-object'.

B.6. Dynamic Geometry.

It is sometimes desirable to have an object that can change shape. The shape might simply be deformed, such a bouncing ball or the squash/stretch used in classical animation, or it might actually undergo metamorphosis into a completely different geometry. The latter effect is commonly known as 'morphing' and has been extensively used in films, commercials and television shows. Morphing can be conducted in the image domain, 2D morph, or in the geometry domain, 3D morph. The latter is applicable to VR systems. The simplest method of doing a 3D morph is to pre-compute the various geometries and step through them as needed. A system with significant processing power can handle real time object morphing.

B.7. Swept Objects & Surface of Revolution.

A common method for creating objects is known as sweeping and surfaces of revolution. These methods use an outline or template curve and a backbone. The template is swept along the backbone creating the object surface, or rotated about a single axis to create a surface of revolution. This method may be used to create either curve surfaces or polygonal objects. For VR applications, the sweeping would most likely be performed during the object modelling phase, and the resulting polygonal object stored for real-time use.

Appendix C

Three-Dimensional Modelling and Visualisation Engines Involved in Preliminary Examinations

Rend386

Rend386 is a free DOS based program which enables VR hobbyists to create Virtual Environments (VEs). This program has proved very popular and many device drivers and converters have been written to allow the use of low cost peripherals and to import file formats such as .DXF. Objects are created as an ASCII file in Rend386's .PLG file format. Multiple spot-lights and point source lights are available to add more realism to the environments created. Simple object texture is also possible using Rend386's preset textures.

VR386

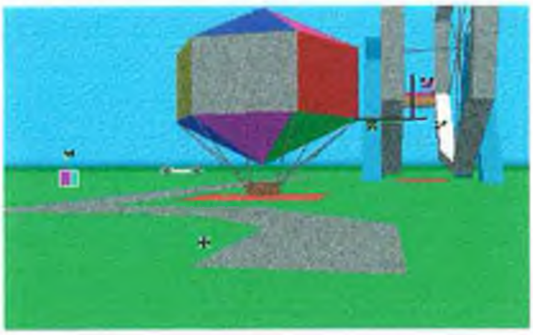

VR386 is the a follow up to REND386 which provides the designer with some additional features including a better time-base for timing and control in the Virtual Environment (VE), the use of extended memory and a more simplified programmer interface.

World Builder.

World Builder is similar in appearance to Rend386, although it does have twice the vertical resolution and extra features, including the use of extended memory. World creation is performed simply using a mouse and drop down menus and an extra packages included to allow the creation of shapes using a CAD packages. World Builder uses the Rend386 .PLG file format and .WLD world files, which allows the use of preset glass and metal textures found in Rend386. A similar scripting language to that used in Rend386 is available, and an interface to C can be used for special requirements.

Avril

In order to provide cross platform support, and make further development simpler, Rend386 has been completely rewritten in standard C to form A Virtual Reality Interface Library (AVRIL). This new package can still use the .PLG objects and .WLD world files but using the new graphics code has made the process much simpler.

VROOM

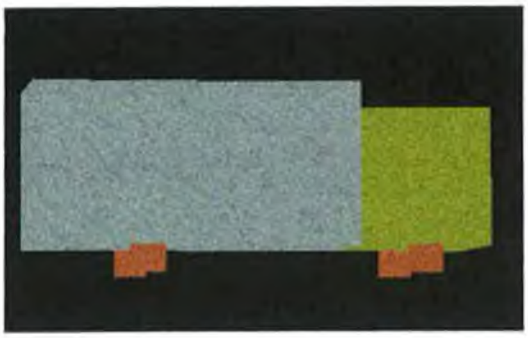

Virtual Reality Using Object Oriented Methods (VROOM) was developed to provide a VR Viewer for the author's SWOOPVR industrial VR system. VROOM uses object oriented techniques to provide simple programming of Virtual Worlds. Creating objects from scratch in code can be an extremely labourious and impractical experience. VROOM has its very own ASCII file format for describing 3D objects, called the .3D format. The geometry of the world is described by VROOM by first describing all the objects' coordinates, and then listing each face by describing its colour, shade and the list of face coordinates in an anti-clockwise fashion. Converters are available to convert from the Rend386 PLG format.

3D-ware

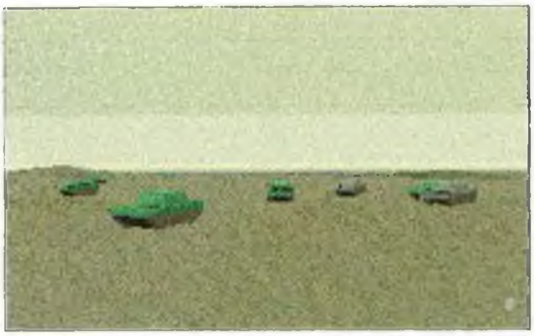

3D-Ware has been written entirely in assembler in an attempt to achieve real-time interaction and also to allow interfacing with a wide variety of high level languages. Many of the tasks such as creating objects and adding colour to them can be carried out entirely using a graphical tools, therefore, reducing the need for a programming language.

BRender

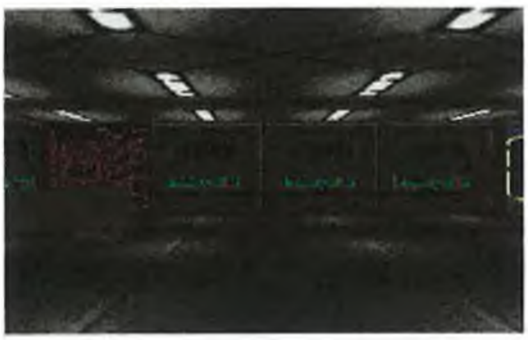

BRender (Blazing Render) is a complete VR system with an applications interface, graphics libraries and device drivers. The program supports gouraud shading, multiple camera and light sources, sprite scaling and collision detection.

World Toolkit

World Toolkit (WTK) is a commercial VR environment manufactured by Sense8 and is available for many platforms from a Silicon Graphics Onyx to an IBM compatible PC. It is intended for single users - although version 2.0 does have some low-level support for networking. The software consists of a library of C-callable functions, and can exchange data with other Windows applications. A separate utility is available which allow Virtual Worlds to be created in WTK using a graphical user interface (GUI) instead of C.

Virtus Walkthrough.

Virtus Walkthrough is a flat shading engine for Macintosh and Windows from the Virtus Corporation. VEs are constructed using a 2D graphical interface, and can be experienced at any time using a 3D window. This approach is used in an attempt to provide a designer with an extremely versatile and intuitive tool for creating VEs.

Appendix D

Three-Dimensional Games Engines Involved in Preliminary Examinations

Wolfenstein 3D.

Wolfenstein 3D, produced by Id Software, was one of the original 3D games engines. Wolfenstein 3D received many awards for its originality, game characteristics and realistic rendering of its environment and characters. Although Wolfenstcin 3D revolutionised the world of 3D gaming it still had limitations if one were to model a real-life simulation. In the Wolfenstein engine, the walls were always perpendicular and in fact were actually four sided blocks with the same wall designs on all four sides. Levels were constructed in Wolfenstein by placing these blocks on a fixed grid. The fixed grid was made up of a 64 x 64 square area. Each square on the grid represented an area approximately 8 ft^2. This made all the rooms constructed seem box-like and confining. The Wolfenstein engine also had fixed doors and ceiling heights that were always 64 units high. Finally, the Wolfenstein engine made very limited use of lighting. Due to these limitations the engine was unacceptable for the construction of all but the very basic Virtual Environments (VEs) and, thus, was not given any further consideration in this research.

Doom.

To understand Doom, one has to take a look at its predecessor, Wolfenstein. The engine used for the adventure game Doom, hereafter referred to as the Id engine, was written by Id based on their experience with Wolfenstein.

Heretic.

The engine used in the game Heretic was an engine produced by the collaboration of Id Software and Raven Software. The engine used in this game was a refinement of the Id Engine. The methods used to construct the worlds were extremely similar to those used in the Id engine. The major advancement of this engine was that it enabled the participant to look up and down and to the front, back, left and right. This allowed the engine five degrees of freedom (5DOF) compared to the Id engines three degrees of freedom (3DOF). The second major advancement on the Id engine was that the Heretic Engine allowed a participant to fly in the Virtual Environment (VE). This gave the participant far more maneuverability to explore the VE. Further exploration was not possible because by the time the full version of Heretic was released a better engine called Hexen, created by Raven software, was available.

Hexen.

As mentioned above, Hexen is another creation by Raven Software. The engine used in the game Hexen, referred to hereafter as the Raven engine, is a refinement on the Heretic engine. The Raven engine includes all the features that Heretic possesses plus some extra features. The major difference between the engines used in both Heretic and DOOM and the Raven engine is the Raven engine's programmability. It allows a designer to add behavioural characteristics with the use of its very own scripting language. This powerful scripting language can be used to create a wide variety effects.

Duke Nukem 3D

The engine in the adventure game Duke Nukem 3D, referred to hereafter as the 3D Realms engine, is a huge progression from those examined above. This engine possesses a few new features such as stereo imaging, high resolution graphics and more realistic behavioural characteristics.

Appendix E

Sector Attributes - The Id Engine

E.1. Floors and Ceilings Detail.

The floor of a sector is the horizontal plane that forms the lower boundary of the sector in space. It is the space on which the participant will walk. The ceiling, on the other hand, is the horizontal space which forms the sector's upper boundary. Both these sectors are truly horizontal and can not slope, thus making the creation of sloped roofs and hills a difficult undertaking. A sector's floor and ceiling each possesses two definable properties: height and textures.

The height value specifies the surface's absolute vertical placement in the range -32768 to +32767 units. Naturally the ceiling is placed higher than the floor and the difference between the two heights provides the apparent height of the sector.

The texture value consists of an eight-character name that tells the Id engine which pattern to render on the screen when showing the appropriate surface.

E.2. Lighting Level.

A sector's brightness level, expressed as a number between 0 and 255, determines how brightly anything occupying the sector will be displayed in the Virtual World.

E.3. Behavioural Characteristics.

The behavioural characteristics of a sector determines whether the Id engine will conduct any special behaviour once the participant enters the sector.

Appendix F

Preliminary Explorations - The Id Engine

This experimentation process started with trying to alter the original Level 1 Mission 1 of the game Doom, illustrated in Figure F.1. At first the designer was only interested in moving vertices and lines and creating simple sectors to examine the effect that this would have on the environment. As is the case in Figure F.2., where only a slight alteration of the original level occurs with just a passage running from the beginning to the end of the level. This one alteration took a lot of time and effort before it could be viewed in a 3D format without the Engine "crashing". Many of the errors that occurred were due to certain attributes of the lines and vertices not being suited to one another.

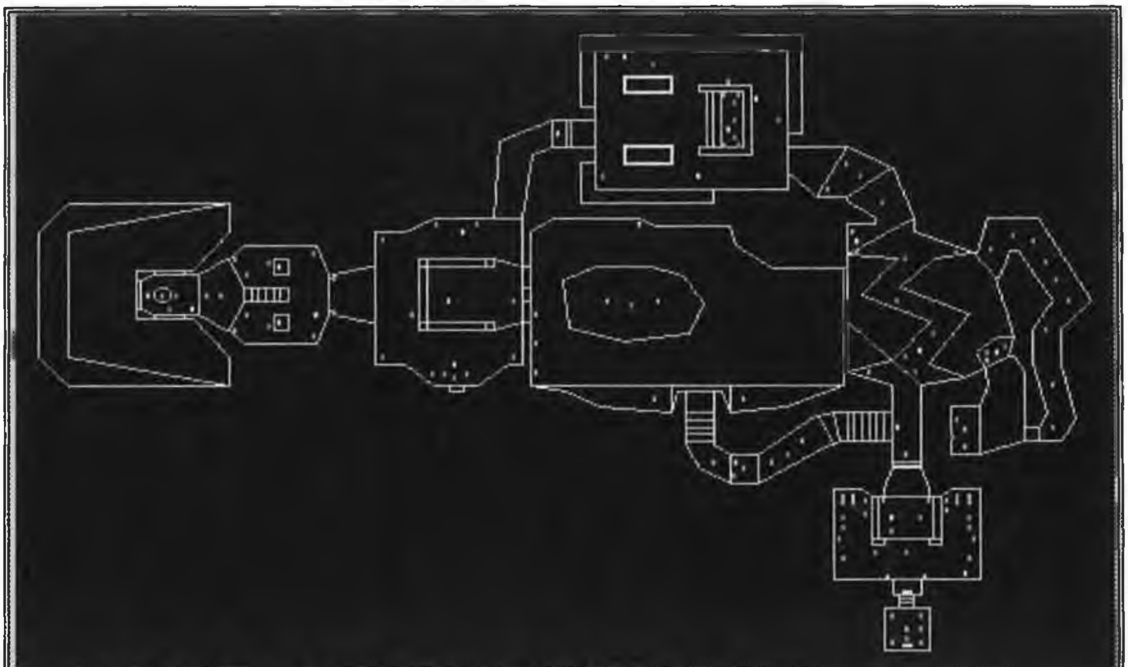

Figure F.1. A Map View of the First Level from the Game "Doom".

The map displayed in Figure F.3. is another progression from the original level and, for the first time, a new room was introduced. This room is linked to the original level by another sector but at this stage to turn this small sector into a door was still unachievable. After the original problems were overcome this level was created quite quickly but to create a door proved to be quite a time consuming problem.

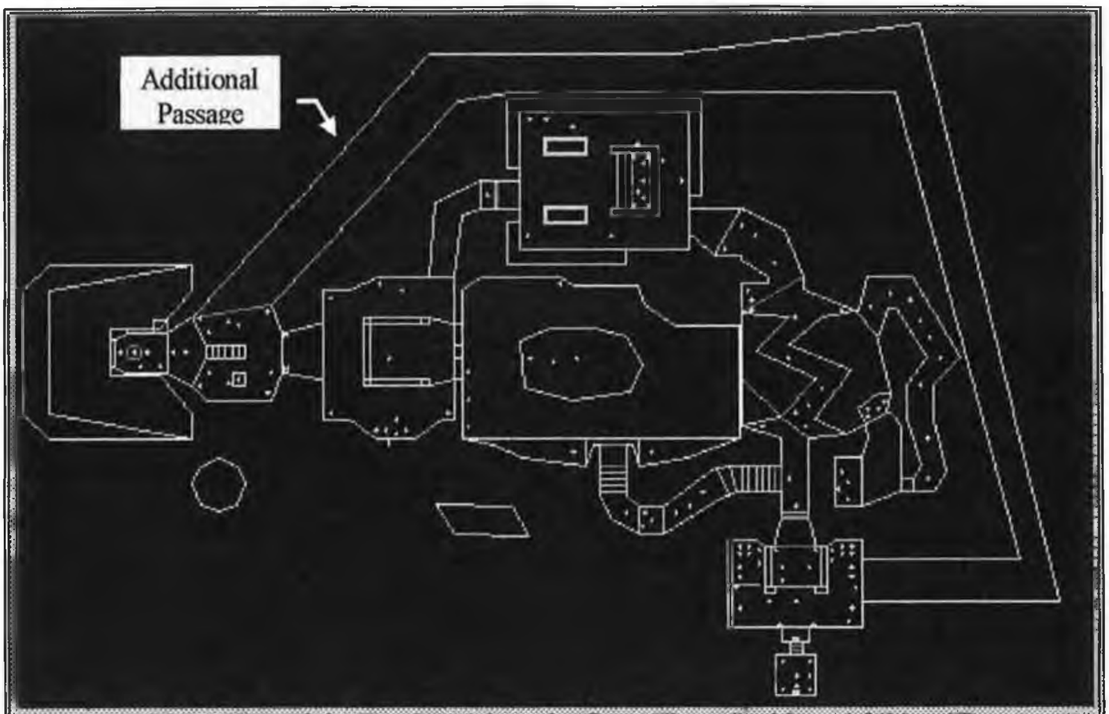

Figure F.2. A Map View of the First Level from the Game "Doom" with an Additional Passage.

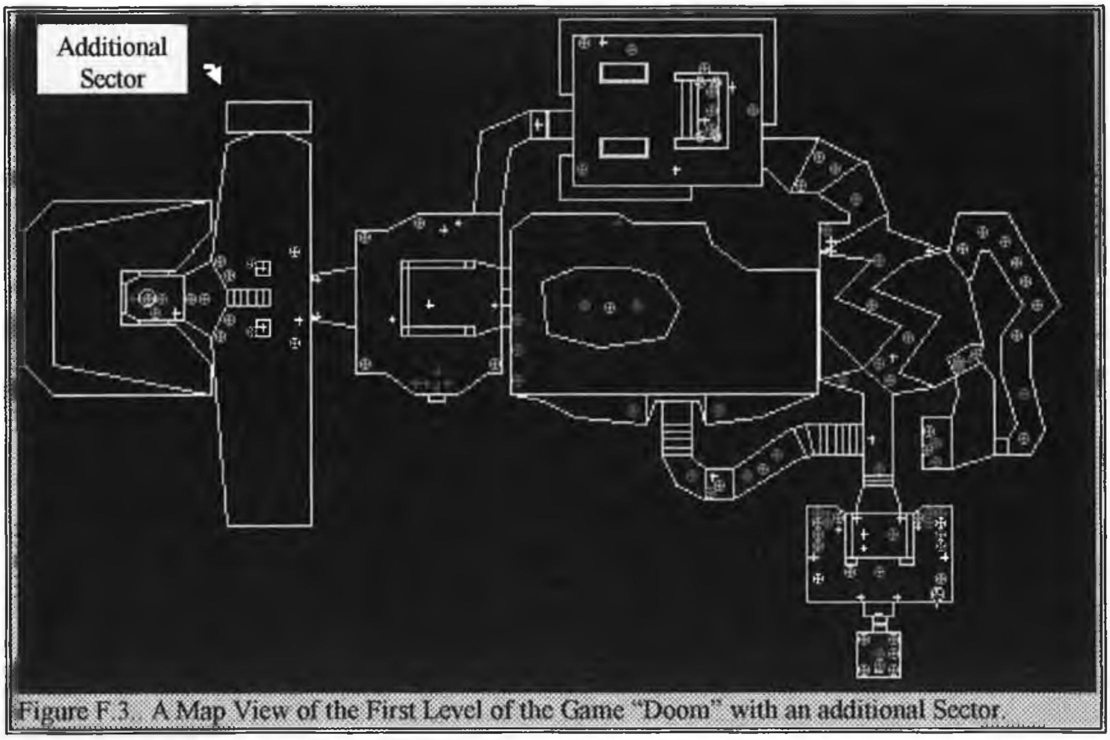

Figure F.3. A Map View of the First Level of the Game "Doom" with an additional Sector.

The map illustrated in Figure F.4. is a bigger level that the level in the previous Figures. This level was the level in which the problem of designing doors was actually overcome. This proved to be a huge milestone in understanding the basic concepts of designing with DEU. At the end of this level a special type of line was placed which, when a participant tripped a switch, the Virtual World ended and the participant moved onto another world.

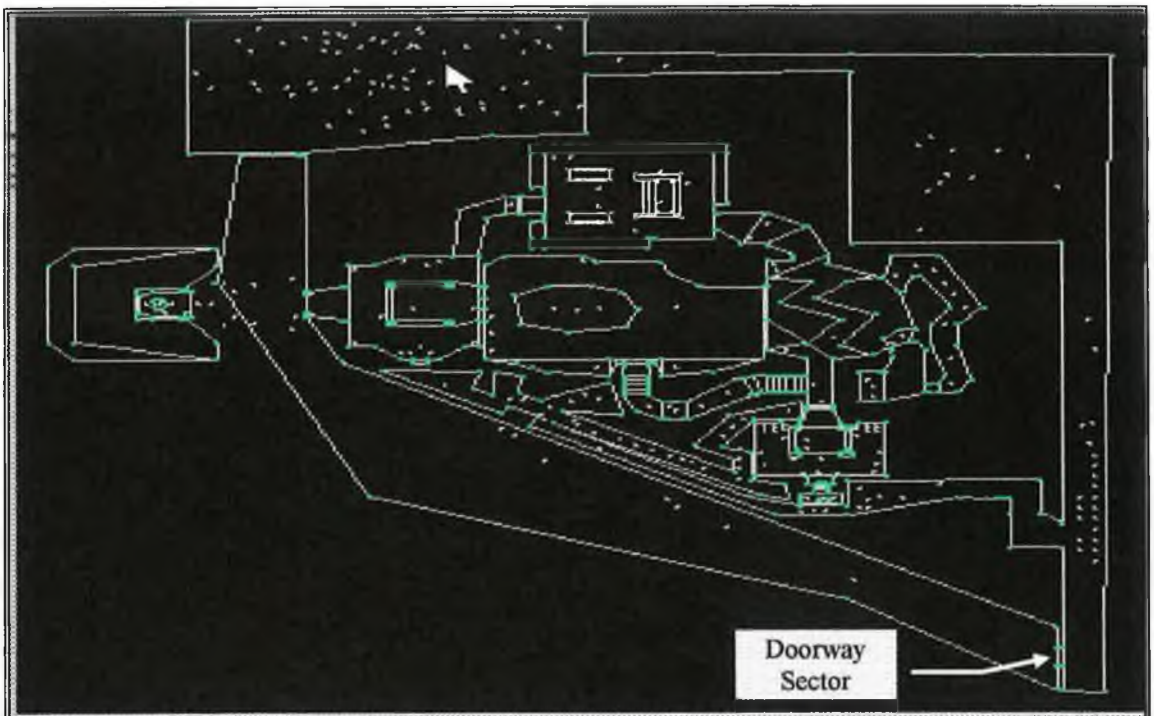

Figure F.4. A Map View of the First Level of the Game "Doom" with an additional Sector.

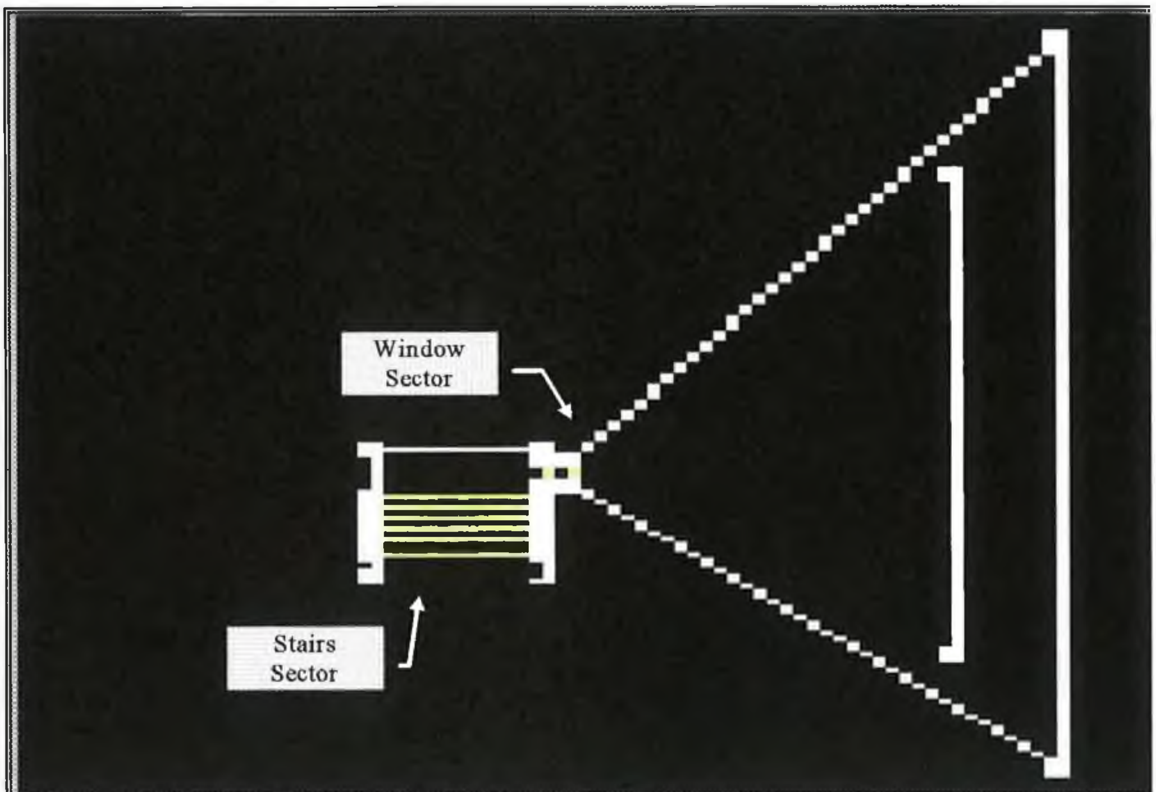

Figure F.5. A Close-up Map View of a Stairs and Window Sector.

In the level above a stairs and a window were placed into it, a close up of both of these appears in Figure F.5. These two additions appeared very easily because most of the problems that would have occurred were overcome in the process that was undertaken to design a door.

Appendix G

Activation Types - The Id Engine

1. Manual activation - Local Activation. Action is triggered by the spacebar. The main difference between this action and the spacebar activation below is simply that this action is a local action. A participant presses a spacebar in order to activate that actual sector, usually a door, but in the spacebar activation a participant presses the spacebar on a certain line in order to activate a remote sector.

2. Spacebar activation - Remote Activation. Actions activated are by pressing the spacebar. These lines are usually used to trip a switch or press a button. A participant must be directly in front of a spacebar activated line in order to have any affect on it. A participant does not have to be right up against one of these lines to activate it, but he has to be within 31 units of a button or switch in order to activate it provided there are no obstacles in the way. When a number of spacebar activation lines are placed close together only the closest line to the participant will be activated once the spacebar is pressed. It is usual to provide a participant with some indication that there is a switch line waiting to be activated. This is normally achieved by placing an appropriate wall texture on the line.

3. Walk-through activation - Remote Activation. These actions are activated as soon as the participant steps over the line. The participant is deemed to have stepped over a line only when his center point makes a transition from one side of it to the other side. This means that there will need to be a gap of at least the participant's radius, 16 units, on each side of any walk-through line for it to ever be activated.

4. Impact activation - Remote Activation. Actions are activated by being shot at. Though all weapons were removed from the game, this special line attribute still remains useful. It was used in the prototype model to activate special sound files to give the participant information about certain rooms but it could also be used to activate many other actions.

Appendix H

Replacing Loading Screens - The 3D Realms Engine

The first piece of information which was replaced was the 3D Realms startup screen, displayed in Figure H.1. In this screen information regarding the version of the engine, the internal and external files, the control device configured and the types of script that were executed by the configuration file is displayed. These vary depending on the different options chosen by a participant using D-VR3. Before participating in a walkthrough of one of the worlds constructed using the 3D Realms engine in D-VR3 a participant is offered a choice of which world they wish to experience and which control device they wish to use during the experience. Once the participant has made their decision regarding these two choices the appropriate world is launched using a configuration file which will load the chosen control device. These tasks are carried out automatically once the participant makes his decision which means that the participant does not need to have to worry about any of these aspects of his experience. For this reason it was decided that the original startup screen, Figure H.1., should be replaced by a more appropriate screen, illustrated in Figure H.2.

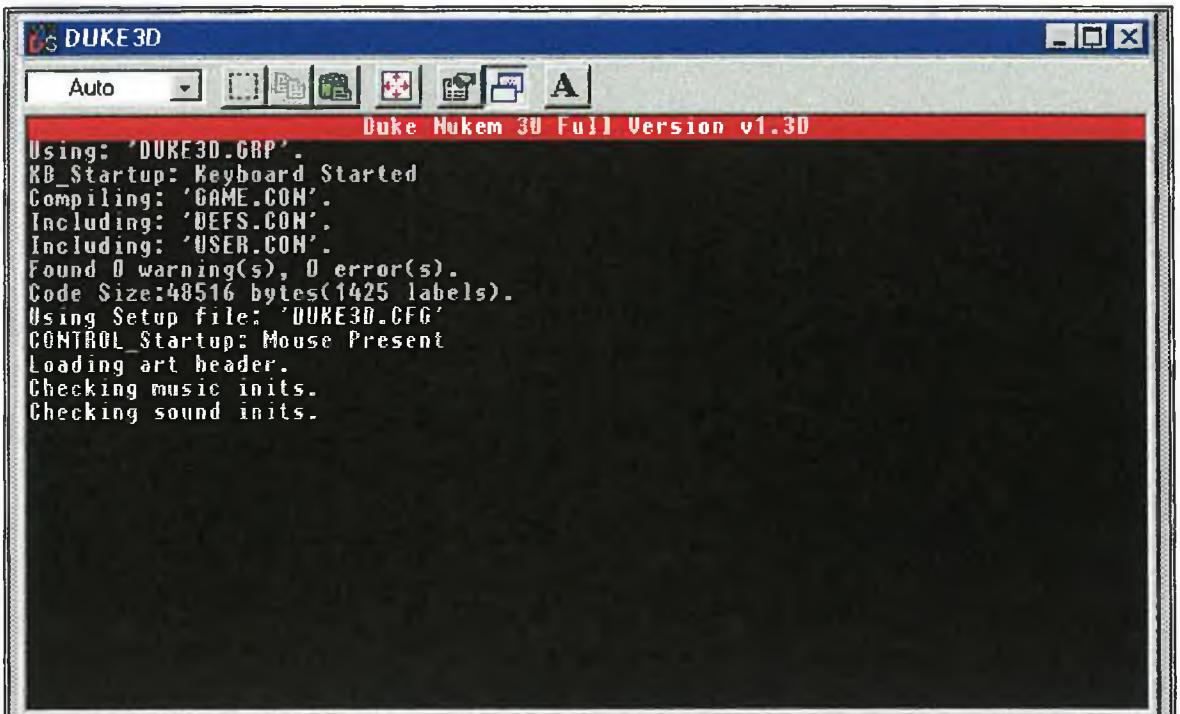

Figure H.1. The Original Startup Screen in The 3D Realms Engine.

The removal of the startup information was an exercise which required a large amount of precision. It involved specific alterations being made to the 3D Realms engine's main executable file DUKE3D.EXE. This process required the use of a hexadecimal editor. The hexadecimal editor that was used in this case was one called Hedit32. HEdit32 is a hexadecimal edit for binary files. This means that it can be used to open any arbitrary binary file in order to edit its contents in either hexadecimal presentation of its data bytes, or in corresponding ASCII characters. Unlike text editors, HEdit32 does not allow a user to insert or delete bytes in the middle of the file, thus, shifting the rest of the file back or forth. Such operations are common mainly for text files, where the absolute location of a character is not critical. Binary files usually rely on a particular offset of a byte from the beginning of the file, and therefore shifting the tail of the file back or forth could make the file corrupt. For this reason this process was a precise exercise. Most of the information in the original startup screen has been removed or, to be exact, has been replaced by spaces, that is except for the information in the title bar. The title bar, Figure H.2. now displays the information that the 3D Realms engine has been modified by the designer.

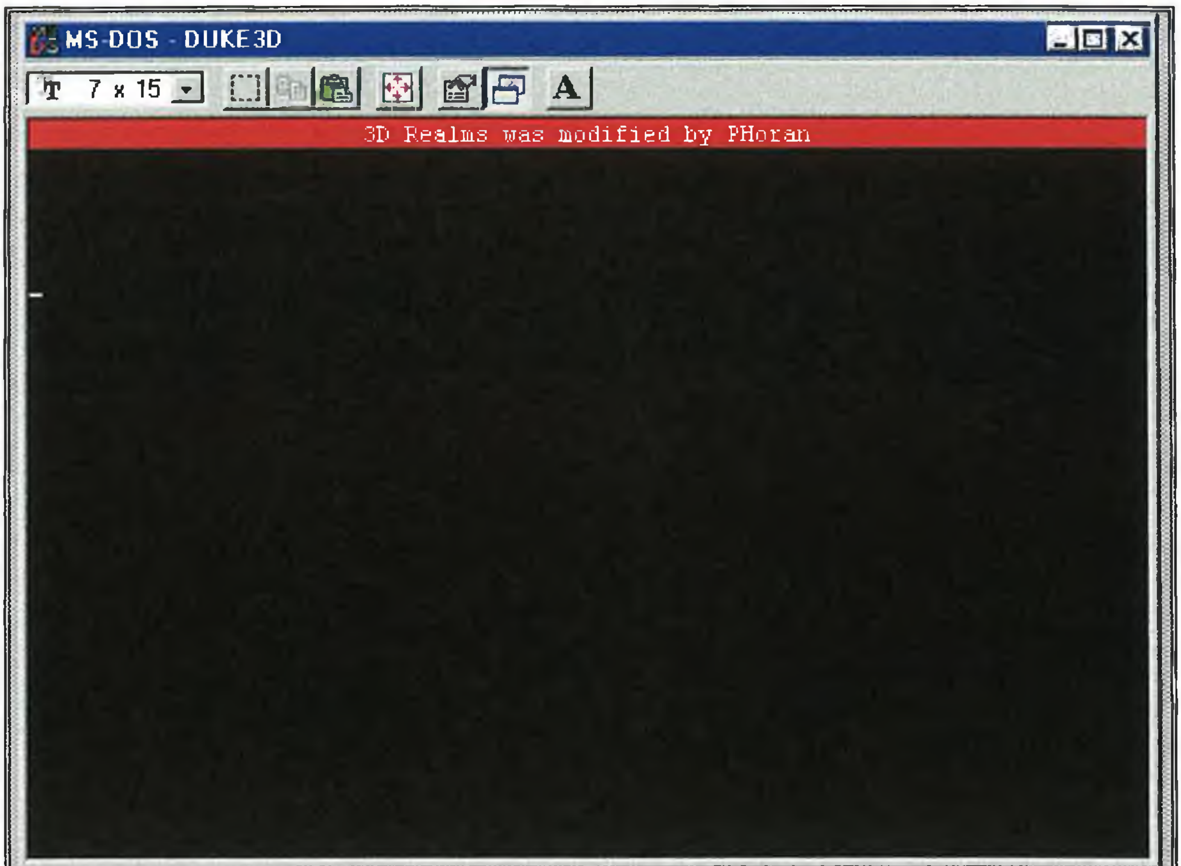

Figure H.2. The Startup Screen in The 3D Realms Engine Used in D-VR3.

Once the appearance of the startup screen had been satisfactorily altered the next task was to alter the screen sequence which followed. The original screen sequence, displayed in Figure H3, was removed. The main reason for removing the original screen sequence was to remove all evidence that this engine was indeed a games engine, which would of course detract from the "feel" of the application.

Figure H.3. The Loading Screen Sequence For The 3D Realms Engine.

Once the starting sequence had been removed successfully the next task was to replace the main menu that a participant would encounter as soon as he entered a Virtual Environment. The original main menu, illustrated in Figure H.4., enabled a participant to choose different game levels and to change some of the game playing options. The first alteration that was made to main menu was that the title of the menu was changed from the "Duke Nukem 3D" to "D-VR3". This alteration involved replacing the original graphic with a similar D-VR3 graphic. The graphic was drawn using Corel Photo-Paint Version 6.0 and was imported into the engine using Editart. After the title had been changed the next step was to customise the options in the menu in order to make them more suitable for D-VR3. These options were changed to enable a participant to chose different maps. Many of the options from the options command were removed and the instructions in the Help command were also changed to suit the worlds that have been constructed. All of these changes were carried out using HEdit32. The final menu is displayed in Figure H.5.

Figure H.4. The Original Main Menu From The 3D Realms Engine.

Figure H.5. The Final Main Menu From The 3D Realms Engine.

The only step that remained in order to make the engine's interface more acceptable for the actual application being developed was to remove weapons held by the participant and to remove the information bar at the bottom of the screen, both of which are displayed in Figure H.6. The removal of the weapons was a two step process. The first step involved the removal of the pistol graphic from the engine with Editart and the second step involved changing the configuration file to disable a participant using the now invisible pistol or changing it to a visible weapon.

Figure H.6. The Original 3D Realms Engine's Game Interface.

The removal of the Information Bar followed the same principle as that for the removal of the pistol. This was again a two step process. This time the first step involved maximising the screen in order to lesson the number of Information Bar symbols displayed and then removing these remaining symbols using Editart. The second step involved altering the configuration file this time to ensure that the screen could not be reduced to show the remaining Information Bar symbols. After removing all these graphics the 3D Realms engine interface was then free from all evidence that it was, in fact, a games engine and was now just a plain screen, as can be seen in Figure H.7., through which a participant could view the Virtual World.

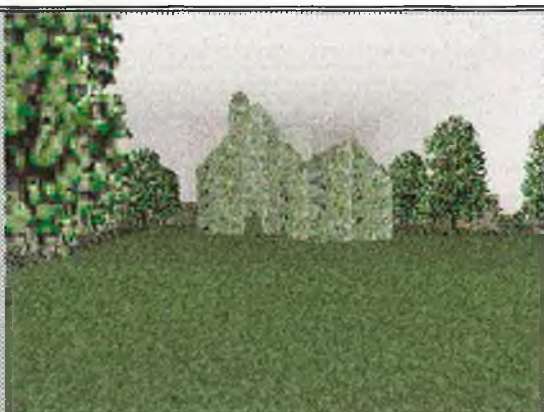

Figure H.7. The 3D Realms Engine's Interface Used In D-VR3.

Appendix I

Other Considerations When Designing a Virtual Environment

I.1. Resolution.

There are different graphics modes in which the computer generally operates. These different modes are usually called the computer's resolution or display mode. The most common resolution on the PC are 640 x 480, 800 x 600, 1024 x 768, and 1280 x 1024. This means, for example with 600 x 480, the picture shown on the screen is made up of 480 lines which are each 640 dots wide. Each one of these dots is referred to as a pixel, which comes from the phrase picture element. These pixels are all the same size. Therefore, in the example of the 640 x 480 resolution, the screen has 307,200 pixels to display a particular image. This may seem like a lot but it is not. In order to display a desired picture on the screen, each pixel is assigned a colour and a brightness value. That value is usually stored on the memory chips that reside on the display card. The number of colours from which a single pixel can choose and the number of colours that the screen can display are decided by the pixel depth. Pixel depth is the number of bits it takes to display one pixel. Eight bits are equal to one byte. A single byte can have 256 possible values.

The major drawback of this mode is that one can quickly run out of colours. Most pictures are made up of thousands, possibly millions different colour shades and hues. With only 256 colours there may not be enough different colours to display the picture in a recognisable fashion. In this mode a colour palette is used. The palette has 256 slots in it, each one with its own unique colour. A single palette slot is called an index. Each index is numbered 0 - 255 and each represents a combination of three colours: red, green and blue. Each colour can have a brightness level from 0 - 63. By mixing the three colour channels in various intensities, almost any colour required can be created. Other display cards can display up 16.7 million colours.

The next step up from 8 bits per pixel is 16 bits per pixel. The main difference between this and the previous mode is that this mode does not use a colour palette. This mode takes 2 bytes of memory to display one pixel, and is sometimes referred to as high colour mode. This mode can yield up to 65,536 simultaneous colours. Although this mode displays significantly more realistic images it requires twice as many bytes to display a picture. This means that twice as much storage space is required and it takes twice as long for the picture to be drawn on the screen. For this reason most computer games and packages which try to achieve real-time interaction still depend on the 256-mode to display their graphics.

The final mode is usually called true-colour mode because, theoretically, it can display all colours, resulting in images which are true to life. In reality, it can not display all the colours of the spectrum, but it can display all the colours visible to the human eye. In 16 million-colour mode, 3 bytes is required to display one pixel. The storage required to store one 640 x 480 image in this mode is 921,600 bytes, almost 1MB.

The 3D Realms engine is capable of 1600 x 1200 colour resolution. This option is not available in the engine's setup file but can be set using the configuration file. The main reason that this resolution is not included in the setup file may be that at this resolution the frame rate slows to an unacceptable speed.

The resolution used for the 3D Realms engine in D-VR3 is 1024 x 768. At this resolution a frame rate of over 26 frames per second was achieved.

The graphics, although they are not an extremely high resolution, are extremely fast. The 3D Realms engine runs in normal mode, 320 x 200, 256 colour graphics mode, or in VESA/SVGA mode, which can incorporate graphics up to 1600 x 1200, 256 graphics mode. What makes the graphics even more believable from a user's point of view is that the 3D Realms engine incorporates real-time, texture-map rendered graphics. As the participant moves around, the scene changes accordingly as the images are generated in real-time they are not just images being loaded from the hard-disk. They have been rendered in real-time complete with realistic texture-maps. This feature of the 3D Realms engine was one of the main reasons for using the engine in this project.

The resolution of the worlds created with Virtus Walkthrough and VRML all depend on the resolution of the monitor being used.

I.2. Stereo Imaging Display.

All three types of engines discussed above can display their worlds in stereo.

The 3D Realms engine enables a participant to view a world created using the engine in 3D with the aid of red/blue glasses or a "Crystal Eyes" HMD. The worlds can be experienced in these modes by altering the engine's configuration file to allow the appropriate images to be displayed. In the course of this dissertation it was decided to base the Virtual Worlds created on commercially available PC's with no additional hardware. Therefore, the worlds were not tested using the Crystal Eyes Display but were, in fact, tested using the red/blue glasses. The glasses provided a less than satisfactory result and were difficult to endure for any long periods of time. For these reasons, these processes were not further developed.

Environments created using 3D graphics and animation packages can be viewed in stereo by using a stereo 3D Flic player. Virtual Environments (VEs) constructed using Virtus Walkthrough can also be viewed in stereo by first of all changing the stereo option in the walk menu and then using the appropriate configuration. Virtus Walkthrough can support stereo pair, rotated stereo pair, over/under display and alternate scan lines. Stereo VRML is possible with the inclusion of specific display hardware such as HMD's supported by VRML 2.0.

I.3. Navigation.

Navigation is an extremely important consideration in each of the packages used. In Virtus Walkthrough a mouse is used to navigate through the Virtual World. Navigation is based around a point at the center of the Virtual World. The speed at which a participant moves in a world is dependent on how far from the center point the mouse pointer is when it is clicked. The further from the center point the faster the participant moves. To move forward, press the left mouse button above this point and press the left mouse button anywhere below this point to move backward. Turning left or right occurs when the left mouse button is clicked on the appropriate side of the center point. If the participant right mouse clicks above the center point it causes him to rise up and a right mouse click below this point causes the participant to sink down. Right mouse clicking to the left or right of the center point causes a participant to slide left or right. In addition, certain keyboard controls can be used to navigate in a world created in Virtus Walkthrough. A participant's view perspective can be changed by holding down the Shift Key while using the left mouse button. This causes a participant's view to tilt sideways.

The navigation mechanism employed by the 3D Realms engine is relatively straightforward. A participant can move throughout the environment with the use of many different control devices such as a keyboard, mouse or joystick. In D-VR3 the participant is offered a choice before he enters the Virtual Environment (VE). Once the participant has chosen an input device and entered the environment he must continue to use this device until he returns to D-VR3. The control devices to be used can also be changed by using the 3D Realms engine's set-up program. In addition the actual movement that each control dictates can also be changed. The default controls in D-VR3 are:

The Mouse Buttons - Both the left and right mouse buttons are set up to conduct the one function. When either of these buttons are pressed the participant moves forward in the direction that he is facing. This function is set up in order to reduce the amount of mouse movement that the participant has to make.

Arrow Keys - The up arrow key moves the participant forward, the down arrow moves the participant backwards and the right and left arrows turn the participant right and left respectively. These can be used in conjunction with, or instead of, a mouse.

Shift Key - Holding down the Shift Key while the participant is moving allows him to move faster.

[J] Key - This key activates Fly-mode which enables the participant to fly. To fly higher a participant must press the [/] key and for a participant to reduce his altitude the [.] must be pressed. The [J] key can be press again to deactivate Fly mode.

[Home] - The [Home] key enables a participant to look up.

[End] - The [End] enables a participant to look down.

[A] - The [A] key enables a participant to jump.

[Z] - The [Z] key enables a participant to crouch.

Navigation in a VRML depends on the VRML browser being used but in many instances navigation in a VRML is quite like the mouse navigation mechanism described in Virtus Walkthrough. Finally, because an animation file created using 3D Studio is basically non-interactive the only controls used to view such a file are start, stop and pause buttons located at the top of the animation.

I.4. Sound.

The integration of sound effects into Virtual Worlds is clearly a compelling component of any environment. Using vision-only systems would not be doing VR justice as the multi-modality of the experience is an essential feature for any applications and, therefore, sound must be an addition to any VR application to Tourism.

The 3D Realms engine has an open architecture that enables a designer to add new (or modify old) levels, characters, graphics and sounds, in fact any aspect of a VE. This is the aspect of the 3D Realms engine which makes it such a powerful tool for the creation of VEs. The addition of customised sound files to the 3D Realms engine was a relatively straightforward matter which included the creation of the sound file in Creative Wave Studio. These files were then named appropriately and placed into the directory which contains the 3D Realms engine's main executable file. These sound files now replace any sound files contained in the main 3D Realms File, DUKE3D.GRP, with that same name. Certain sound effects were placed in suitable locations in the worlds so that when the locations are passed the appropriate sound file is played. Sound files could contain anything from environmental sounds, such as wind, rain or water flowing, to voice-overs offering the participant more information regarding their specific location within the Virtual World.

Sounds were added to an animation in the form of a voice over and sound effects by simply recording the sound file in Creative Wave Studio and adding the sound file to the animation script. Neither Virtus Walkthrough Pro 2.6. nor VRML Version 1.0. supports sound. However, VRML version 2.0. does support sound.

I.5. Multi-users.

Virtual interface technology has been viewed from its beginnings as a genera-purpose simulation tool. To this point, the majority of its documented applications have been simulations of some type. But these simulations deal with very few virtual objects and rarely have more than one simultaneous user.

The 3D Realms engine allows networking for up to eight participants. This means that up to eight participants can enter and interact with the one environment at the one time by using a serial cable connected between two machines, a modem or a network. These participants can be located in the one room, or on different sides of the world. These participants can see all of the other participants and can communicate with them by pre-determined dialogue or by pressing the [T] key and typing a specific message and then pressing the [Return] key. Participants within the environment could be tour guides or other potential Tourists. Therefore, a participant could not only learn from the surrounding VE but also from other participants they might meet within that environment. This feature has great potential when applied to the Tourism Industry.

None of the other packages evaluated, Virtus Walkthrough, 3D Studio, nor VRML Version 1.0. allow multi-user access.

I.6. Pre-Programmed Walkthroughs.

A Pre-programmed walkthrough of the VE can be recorded on every package used but it is the aim of this research to provide a participant with a live, interactive Virtual World in which they are allowed the freedom to choose which direction to go and which object to interact with. Therefore, although pre-programmed walkthroughs may be very useful in other applications it was not given any further consideration in D-VR3.

Table I.1. The Packages Used to Create Virtual Environments in D-VR3 & Their File Formats.

Packages	File Formats.
Virtus Walkthrough 1.0	.WLK, .VVR.
Virtus Walkthrough Pro 2.0., 2.5 & 2.6 Beta.	.WTP, .WLK, .VVR.
Doom	.Wad (Doom 1 Wad)
Doom2	.Wad (Doom 2 Wad)
Hexen	.Wad (Hexen Wad)
Duke Nukem 3D	.Map
3D Studio	.DXF
Virtual Reality Modelling Language (VRML)	.WRL

I.7. Conversion Standards.

Tools used for the creation of VEs are becoming increasingly more mature simply because there is high demand for such tools in traditional markets, primarily CAD and animation. Early work by both VPL and Autodesk in these areas, combined with Autodesk's dominant position in the desktop CAD market, has meant that 3D DXF is seen as the standard format for creation and transferring of model information between VR systems. This means that the actual construction of VEs can, in many instances, be carried out in any number of conventional computerised-design sites using standard software systems, and the imported into another system. Therefore, the choice of 3D DXF as a standard was adopted in order to design some of VEs used in D-VR3. Figure I.1. in conjunction with Table I.1. illustrates the interaction between the different file formats used in D-VR3. As can be seen a world constructed in Virtus can be translated successfully into .DXF which in turn can be translated to 3D Studio and VRML. A world created in the Id Engine can be translated to any of the other games engines examined as well as VRML. As mentioned in section 5.15.2. VRML Version 1.0. files can be converted to VRML Version 2.0. files.

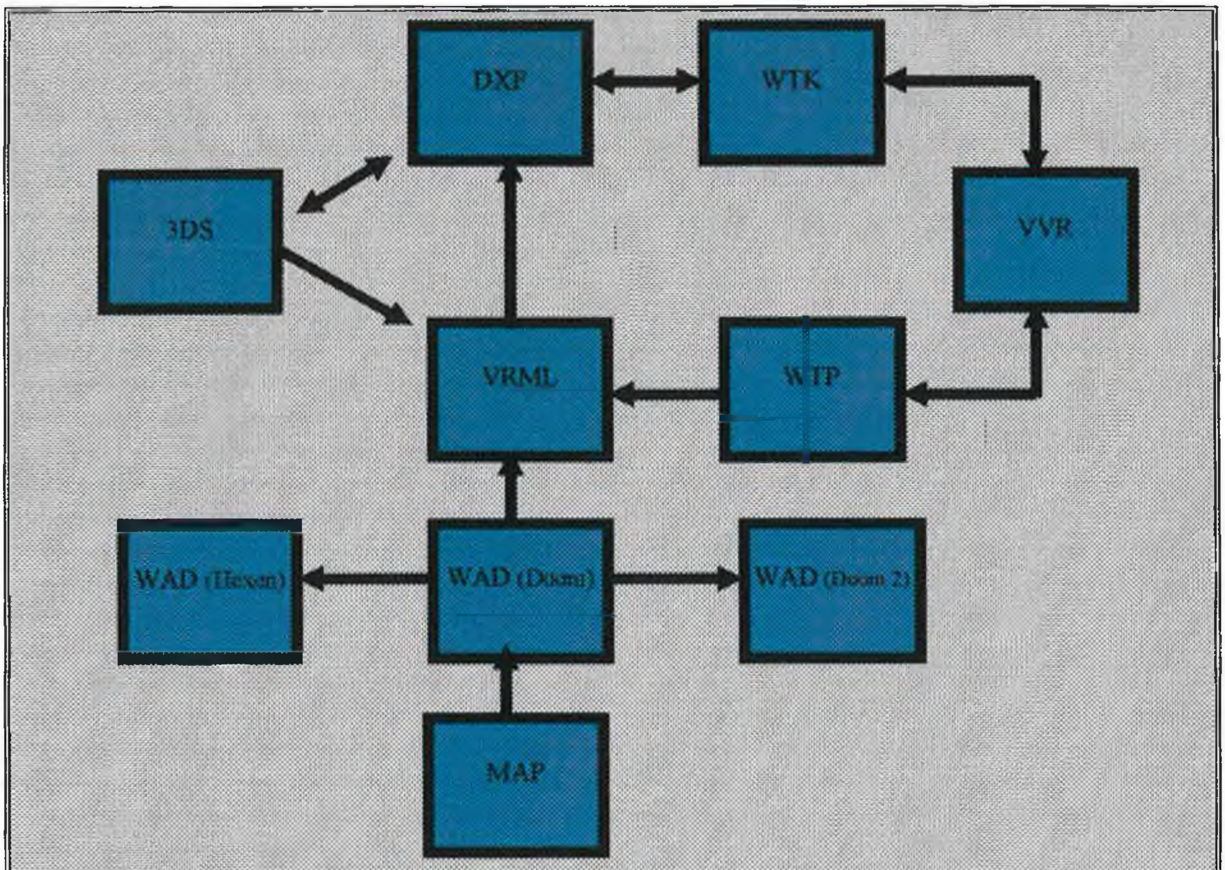

Figure I.1. Standard File Format.

I.8. World Authoring versus Playback.

A Virtual World can be created, modified and experienced. Some VR systems may not distinguish between the creation and experiencing aspects. However, there is currently a much larger body of experience to draw upon for designing the world from the outside. This method may use techniques borrowed from architectural and other forms of CAD systems. Also the current technologies for immersive VR systems are fairly limiting in resolution and latency.

All the VR packages examined had both an authoring mode and a playback mode. The authoring mode may be a standard text editor and compiler system, or it may include 3D graphic and other tools. Such a split mode system makes it easier to create a stand-alone application that can be delivered as a product.

Appendix J

Weighting of Criteria in Internal Evaluation

Weightings Associated With The Sensory Vividness Criteria.

Number of Sensory Channels	Weight
Sight	5
Sound	5
Smell	5
Taste	5
Touch	5
Total	75
Geometry	**Weight**
Scale	5
Scaleable Sprites	3
Anti Aliasing	4
Level Of Detail	4
Animated Skies	2
Sector Over Sector	4
Slopes	5
Portals	4
Animated Furniture	3
Arches	4
Curved Surfaces	4
Total	58
Light	**Weight**
Ambient	5
Directional	4
Multiple	4
Show Cone	3
Hot Spot	2
Shadows	3
Coloured Lighting	3
Projector Effect	2
Total	26
Shade	**Weight**
Wireframe	5
Flat	5
Gouraud	4
Phong	4
Reflective (Chrome)	3
Refractive (Glass)	3
Transparent (Coloured Glass)	3
Total	27

Textures	Weight
Plain	5
Library	4
Real-life Import	5
Transparent	4
Alignment	4
Shading	3
Animated Textures	3
Mapping	4
Bump Map	4
Total	36
Total Sensory Vividness	222

Weightings Associated With The Interactivity Criteria.

Input Devices - Transducers.	Weight
Mouse 2D	5
Keyboard	5
Joystick	5
Mouse 3D	3
Trackball	3
Spaceball	3
Touchscreen	3
Wand	3
Sensors	3
Gloves	3
Bodysuits	2
Retinal Scanners	2
Speech Transducers	2
Thought Transducers	2
Total	43
Output Devices - Displays.	**Weight**
Computer Monitor	5
Head Mounted Display	4
BOOM	4
Red/Blue Glasses	3
Shutter Glasses	3
Cyberscope	3
Projected Image	3
Sound Device	4
Stereo Sound Display	4
Haptic Displays	4
Modems	3
Network	3
Serial/Parallel	3
Total	46

Range of Interactivity Offered	Weight
Look at, but not touch, objects	5
Interact with objects	5
Interact with and change objects	5
Total	15
Degrees Of Freedom	**Weight**
Left/ Right	5
Forward/Backward	5
Up/Down	5
Pitch	5
Roll	5
Yaw	5
Total	30
Specific Interactivity Features	**Weight**
Sliding Doors	4
Swinging Doors	4
Mirrors	3
Moving Sectors	3
Earth Movements	3
Demo Cameras	3
Breakable Glass	3
Echo	3
Total	26
Scripting	**Weight**
Motion Scripts	4
Trigger Scripts	4
Connection Scripts	4
Procedural Modelling	4
Simple Animation	4
Interaction Feedback	4
Total	24
Total Interactivity	185

Weightings Associated With The Time Criteria.

Time	Weight
Update Rate	5
Latency	5
Total	50
Total Time	50

Appendix K

The Questionnaire

Questionnaire.

1. Do you currently have access to a computer? *(Please Tick One)*

   Yes ❑
   No ❑

2. If the answer is Yes, is this access *(Please Tick One or More)*

   At Home ❑
   At Work ❑
   At College ❑
   Other *(Please Specify)* ____________________ ❑

3. Would you be likely to use this program as a form of Tourism Information? (*Please Tick One)*

   Yes ❑
   No ❑

4. If the answer is Yes, would you do so *(Please Tick One or More)*

   prior to your visit ❑
   during your visit ❑
   after your visit ❑

5. Why would you be likely to use such a program as a form of Tourism Information?______

   ____
   ____

6. Where would you be likely to use such a program? *(Please Tick One or More)*

   Home ❑
   Work ❑
   Tourist Office ❑
   Internet ❑
   Hotel ❑
   At the Tourist Site ❑
   Information Bureau ❑
   Other *(Please Specify)*____________________ ❑

7. Would this type of program make you *(Please Tick One)*

   More likely to visit a location ❑
   Equally likely to visit ❑
   Less likely to visit ❑

8. How do you feel that this program compares to traditional types of Tourist Information (i.e. Tourist Offices, Brochures, Videos, Etc.) *(Please Tick One)*

   Much Worse ❑
   Slightly Worse ❑
   Equally Well ❑
   Slightly Better ❑
   Much Better ❑

9. Have you visited the Monastic Enclosure at Glendalough? *(Please Tick One)*

   Yes ❑
   No ❑

10. If the answer is Yes, how do you feel that this program compares to the actual building. (*Please Tick One)*

- Much Worse ❑
- Slightly Worse ❑
- Equally Well ❑
- Slightly Better ❑
- Much Better ❑

11. Which model would you be more likely to use *(Both models will be indicated during the presentation)*

- First Model. ❑
- Second Model ❑

12. Why?____

____

____

13. Gender *(Please Tick One)* Male ❑ Female ❑

14. Age Group *(Please Tick One)* Below19 ❑ 20-29 ❑ 30-39 ❑ 40-49 ❑ 50-59 ❑ Above 60 ❑

15. Tourist Type *(Please Tick One)* Individual ❑ Tour Group ❑ Family ❑ Day Outing ❑ Other *(Please Specify)* ____________________ ❑

16. Occupation:____

17. Nationality:__________________________

18. Any additional comments would be much appreciated______________________

____

____

____

____

____

____

____

____

____

____

____

____

____

____

____

____

____

____

The Author would like to extend his gratitude to everybody who took the time to complete the questionnaire.

Appendix L

The Specific Aims of Each Question in The Questionnaire

Question 1. This question was used to determine the percentage of the respondents who, at the time of the administration of the questionnaire, had access to a computer. A dichotomous question was used so as to allow for a quick response from those being questioned, thereby getting the questionnaire off to a favourable start.

Question 2. This is a filter style question as only those who responded positively to the previous question were requested to answer it. The aim of this question was to establish where the respondents had access to a computer. This question is a polymorphous question as it is a combination of a multiple-choice question and an open-ended question. If the respondent had access to a computer at any other source bar those displayed in the question they were asked to specify the source where the computer could be accessed.

Question 3. This question was designed in order to determine how many of the respondents would be likely to use D-VR3, or a similar application, as a form of Tourism information. This information was sought with the aid of a dichotomous style question.

Question 4. This question also employs a filter style approach as only those respondents who responded positively to the previous question were requested to answer this question. The aim of this question was to determine whether respondents would use D-VR3 or a similar application prior to their visit, during their visit, or after their visit to a particular Tourism location. A multiple-choice format was used to gain the relevant information.

Question 5. This question was designed in order to determine why each respondent would be likely to use D-VR3, or a similar application, as a form of Tourism information. This question was designed in order to allow each respondent to state what it was they found useful about such an application that would help them in making a decision in regards to a Tourism location. This question necessitates an open-ended format to allow the respondent complete flexibility in their answers.

Question 6. This question is classified as a polymorphous question as it is a combination of a multiple choice format and an open-ended question. The aim of this question was to establish where the respondent would be likely to use D-VR3, if such an application were available.

Question 7. The main purpose of this question was to determine whether the use of D-VR3 would make the respondent more likely, equally likely or less likely to visit a location. This information was sought using a multiple-choice format.

Question 8. The aim of Question 8 was to determine how each respondent felt that D-VR3 compared to the traditional types of Tourism information. A multiple choice format was adopted in the form of a Likert scale ranging from "Much Worse" to "Much Better" in an attempt to gain the necessary information.

Question 9. This question was a basic dichotomous type question asking whether the respondent had visited the Monastic Enclosure at Glendalough. This question may seem a little peculiar considering the questionnaires were administered at the Interpretive Center at Glendalough but many people visit the Center prior to visiting the actual Monastic Enclosure.

Question 10. This question uses a filter style approach whereby only respondents who responded positively to the previous question were requested to answer it. This question, using a multiple choice format, aimed to determine how the respondents felt that D-VR3 compared to the actual buildings in the Monastic Enclosure. The responses were measured on a Likert scale ranging from "Much Worse" to "Much Better".

Question 11. The aim of this question was to establish which of the models indicated during the presentation would the respondent be more likely to use. This question employed a multiple-choice format.

Question 12. This is a filter style question requiring the respondent to identify their reasons for choosing the model they have chosen in Question 11. An open-ended format was used in order to allow the respondent amply scope to describe these reasons.

Question 13. This question is a multiple choice question which is aimed at determining the gender of each of the respondents.

Question 14. The aim of this question is to determine into which age group each respondent's belongs. In order to achieve this aim a multiple-choice format was employed.

Question 15. This question is a polymorphous question as it is a combination of a multiple-choice format and an open-ended question. The main purpose of this question was to identify into which Tourist type each respondent belongs.

Question 16. This question is aimed at establishing the occupation of each of the respondents. The question necessitates an open-ended format to allow the respondent complete flexibility in their answer.

Question 17. This question, employing an open-ended format, was used to determine the nationality of each of the respondents.

Question 18. The final question was designed in order to allow respondents to make any additional comments about the presentation, application or questionnaire. An open-ended format was employed.

Appendix M

Cross-Tabulation 1

Respondents Who Would Use D-VR3 at Home

BY

Respondents Who Have Access to a Computer at Home

Respondents Who Would Use D-VR3 At Home BY Respondents Who Had Access to D-VR3 At Home.				
		Respondents Who Had Access to D-VR3 At Home		
	Count	No	Yes	Row Total
Respondents Who Would Use D-VR3 At Home	No	86	48	134 62.0%
	Yes	14	68	82 38.0%
	Column Total	100 46.3%	116 53.7%	216 100%

Appendix N

Cross-Tabulation 2

Respondents Who Would Use D-VR3 at Work

BY

Respondents Who Have Access to a Computer at Work

Respondents Who Would Use D-VR3 At Work BY Respondents Who Had Access to D-VR3 At Work.

		Respondents Who Had Access to D-VR3 At Work		
	Count	No	Yes	Row Total
Respondents Who Would Use D-VR3 At Work	No	122	56	178 82.4%
	Yes	14	24	38 17.6%
	Column Total	136 63.0%	80 37.0%	216 100%

Appendix O

Cross-Tabulation 3

Respondents Who Would Use D-VR3 at Tourist Office

BY

Nationality

Respondents Who Would Use D-VR3 At A Tourist Office BY Nationality (Part 1 of 2)

Respondents Who Had Access to D-VR3 At Work

	Count	Irish	French	Greek	American	Austrian	Row Total
Respondents Who Would Use D-VR3 At A Tourist Office	No	22			40	4	106 49.1%
	Yes	46	6	2	20	2	110 50.9%
	Column Total	68 31.5%	6 2.8%	2 0.9%	60 27.8%	6 2.8%	216 100%

Respondents Who Would Use D-VR3 At A Tourist Office BY Nationality (Part 2 of 2).

Respondents Who Had Access to D-VR3 At Work

	Count	English	Australian	Swedish	German	Swiss	Row Total
Respondents Who Would Use D-VR3 At A Tourist Office	No	16	4	10	4	6	106 49.1%
	Yes	8	6	16		4	110 50.9%
	Column Total	24 11.1%	10 4.6%	26 12.0%	4 1.9%	10 4.6%	216 100%

Appendix P

Cross-Tabulation 4

Why Respondents Would be More Likely to Use a Certain Model

BY

Which Model Each Respondent Would be More Likely to Use

Why Respondents Would be More Likely to Use a Certain Model By Model (Part 1 of 2)

Why Respondents Would be More Likely to Use a Certain Model

	Count	N/A	Higher Quality	Faster	User Friendly	Presence	Row Total
Models	N/A	28					28 13%
	Virtus Walkthrough		54	2	4	4	100 46.3%
	3D Realms Engine	2	6	22	38	12	88 40.7%
	Column Total	30 13.9%	60 27.8%	24 11.1%	42 19.4%	16 7.4%	216 100%

Why Respondents Would be More Likely to Use a Certain Model By Model (Part 1 of 2)

Why Respondents Would be More Likely to Use a Certain Model

	Count	More Realistic	More Time To Examine	Less Loading Time	Row Total
Models	N/A				28 13%
	Virtus Walkthrough	30	6		100 46.3%
	3D Realms Engine	4	2	2	88 40.7%
	Column Total	34 15.7%	8 3.7%	2 0.9%	216 100%

Appendix Q

Multi-Participant Capabilities of The Environments Evaluated

Multi-Participant Capabilities of The Environments Evaluated			
	Modem	IPX	Serial/Parallel
Id Engine	4	2	2
Raven Engine	8	2	2
3D Realms Engine	8	2	2
Virtus Walkthrough	N/A	N/A	N/A
VRML	Depends on Server		