

VRML, Virtual Reality and Visualisation: The best tool for the job?

G.T. Goodrick

Abstract

Recent developments in Virtual Reality and the techniques of Visualisation, coupled with the growing availability to archaeologists of increasingly powerful hardware configurations, has led to the possibility of engaging with archaeological data in a much more experiential way. The parallel development of the internet has served to provide a distribution medium for this data. What has been lacking until now however, has been a way of combining the two. The missing link that has emerged, enabling us to harness the synergy created between the potentials of both Visualisation and the internet has come in the form of VRML, a technique which is already being actively used to represent archaeological data (Gillings and Goodrick, 1996) and in education (Sander and Gay, 1996). This paper discusses the advantages and also the shortcomings of VRML. The relative merits of VRML and other virtual reality systems are investigated, and an overview of those currently available is presented. It is thus hoped to show how a problem rather than methodological driven approach is required if we are to make best use of the developing technologies. The emphasis throughout will be upon methodologies available to archaeologists themselves, rather than "concept" projects developed by computing companies.

1 Introduction

Since its conception in 1994, the Virtual reality Modelling Language (VRML) extension to the World Wide Web (WWW) has become the standard distributed virtual reality system. It has also done more than anything else to raise archaeological awareness to the fact that building virtual environments is a feasible archaeological technique for the investigation of both ideas and primary data. VRML browsers are now included as standard with WWW browsers, and simple-to-use tools for the generation of VRML worlds are readily and cheaply available. Most peoples hands on experience of Virtual Reality will therefore come as a direct result of using a VRML system. One resultant danger however, is that VRML will come to be seen as being synonymous with Virtual Reality.

This paper attempts to draw upon experience gained over the last 3 years in the design and implementation of virtual worlds in both VRML and other more proprietary systems. The aims of the discussion are twofold: to show firstly that rather than contributing to a generic Virtual Reality, different systems produce models that are suitable for very specific purposes, and secondly to highlight some of the pitfalls encountered during this pioneering work. It should be borne in mind throughout this discussion that the research has been aimed towards low-end systems. As a result all of the capabilities and

limitations discussed relate to standard desktop systems running affordable software.

2 The role of VRML

The practical application of the first release of VRML (henceforth referred to as VRML 1) to produce interactive archaeological representations, including a number of worked case-studies, has been described in detail elsewhere (Gillings and Goodrick 1996). Since the publication of this account VRML 1 has been superseded by VRML 2 which, along with the closely related VRML 97, has become a standard specification (Bell, Carey and Marrin 1996). As this has had major implications for the exploitation of VRML, I would like to begin by highlighting the difference between the two specifications, in particular relation to archaeology.

The primary goal in the development of VRML 2 was the addition of interaction and animation to a previously static specification. Although these facilities did exist in some VRML 1 browsers, they were implemented in a variety of different ways and therefore could not be utilised safely or reliably. The ability to produce both visually appealing and amusing models is therefore enhanced, but is of little direct relevance to archaeologists, who will not gain significantly when the considerable efforts involved in defining movement and behaviour are taken into consideration. One obvious exception would be in the

development of educational materials (Sander and Gay 1996).

A number of other enhancements to the specifications which were added at the same time are, however, much more likely to prove useful to the archaeologist. Before going on to discuss these additions it is important to consider briefly the principal aim of any model building exercise. The primary aim in using VR in archaeology should be to provide a basis upon which the viewer can and will project and visualise a reality (Gombrich, 1960). The construction of a 'perfect' representation of a reality is neither possible or often desirable as discussed in detail in Gillings and Goodrick (1996) and Gillings (this volume). However the ability to provide a data rich environment does enhance the capacity of a virtual reality system to deliver pertinent information and facilitate narratives. For instance, it is not possible to create a snow covered glade through which you can walk, feeling the cold wind from the west, smelling the pine trees in the distance, hearing the snow crunch under your feet as you feel it compress, and watching as the snow scintillates light in millions of directions or your footprints slowly turn to slush. We can however use a number of visual and audible clues to communicate the concept of a snow covered glade and provide information about that environment. For instance by adding a low-fi hubbub of sound, growing in volume with increased proximity, we can indicate the presence of a nearby settlement. A number of enhancements included in the VRML 2 specification make the construction of much richer environments possible and the most significant of these additions will now be discussed.

3 From VRML 1 to VRML 2: enhancements and additions.

The most obvious additions to the specification comprise sound, collision detection, backdrops and scripting support. Both point and directional sound is now supported in the specifications and the presence of collision detection makes virtual worlds more believable, as objects can now be defined as truly solid. As a result we will hopefully no longer be able to find ourselves lost *inside* the fabric of Hadrian's Wall. The use of backdrops and fog effects enable more experiential visibility effects such as clarity to be incorporated within constructed models. One major enhancement of particular interest to landscape archaeologists is the addition of a height field for the efficient representation of topography. This is both easier to use and modify, faster to display and far more compact than the indexed face set

representations utilised in VRML 1. It should be noted however that this feature has not yet been consistently implemented across browsers.

The addition of support for scripting languages and prototyping makes it possible to customise and extend the capabilities of VRML2. We can develop standard libraries of archaeological feature prototypes and define behaviours based on scripts. As a result it would be possible to produce a settlement based on standard buildings which could grow and change over time based on its own internal logic.

Although most of these features were available in various forms in VRML 1 browsers, they were not supported in the formal specification and were therefore not consistently implemented. In consequence VRML 2 can be seen as a significant improvement over VRML 1. There are still however a significant number of problems and limitations with VRML as it stands.

4 Problems with VRML.

The principal limitations with VRML at present can be broken down into 2 broad areas concerning issues of standardisation and consistency, and capabilities respectively. Each will be discussed in turn.

4.1 Issues of consistency

A number of internal inconsistencies exist within the standard and between implementations of the standard provided by different browsers. For instance, both colour mapping and lighting is interpreted differently by different browsers, the same browser on different platforms and even different releases of the same browser (figure 1). We do not have a consistent colour/lighting model and it is therefore very difficult to predict how worlds will look on different systems, especially, given the speed of development, in the next generation of browsers. This situation is complicated further still by the fact that mip-mapping and bi or tri-linear filtering (methods of improving the appearance of textured surfaces) are optional. Support for scripting is as yet optional and as such difficult to incorporate into distributed models. Furthermore a number of features, for example height fields and fog, are only implemented in a subset of browsers. These inconsistencies greatly undermine VRML's goal of becoming a standard distributed virtual reality modelling system and as a result many Internet content providers have to host different versions of

VRML worlds to run on different systems, far more will not host VRML at all.




Figure 1: The same scene rendered using a variety of browsers on Intel (PC) and Silicon Graphics (SG) platforms to show variations in texture mapping. Note: All the images were taken from the same viewpoint, variations in the position of the horizon are due to other browser inconsistencies.

VRML compatibility is a moving goal. New browsers are released frequently and are often only released in a time limited form. This often has the effect of rendering previously constructed worlds unusable. The VRML worlds used in Gillings and Goodrick(1996) had to be modified between submission and publication, and again only a few months later browsers were unavailable to display the worlds as the authors intended. Furthermore at least one more version of the VRML specifications will be released to include a binary format. Just as VRML 1 and 2 are largely incompatible we have no reason to assume any more compatibility between VRML 2 and 3.

4.2 The issue of capabilities

The capabilities provided by VRML 2 still fall far short of those desired for the production of virtual worlds. Texture mapping is provided but is usually limited to a crude resolution of 128 by 128 pixels. Although larger texture maps have recently become available on some browsers the performance overhead makes them all but unusable. This is apparent in, for example, the virtual tomb of Menna (Mitchell, 1996), where textures are reduced on standard PC systems to the point where the content of the images is almost totally lost. Sound is currently supported but acoustic properties are not, therefore sound cannot be realistically represented, and may even contradict visual clues. Sonification (the

production of sound from abstract data) is not supported therefore ruling out a number of potentially rewarding experiential data representation techniques. These limitations described in more detail in Nadeau (1997).

A number of problems arise when data destined to be utilised in a VRML world originates in a CAD or GIS package. VRML's use of a screen based co-ordinate system is inconsistent with real world co-ordinate systems. A Transformation Matrix provides a limited fix but has a number of knock-on implications, such as the inability to incorporate gravity effects and the introduction of a further level of complexity into the process of editing the VRML so produced. The other alternative is to swap the x and z co-ordinates, but the software needed to routinely accomplish this process is unavailable for all but the simplest of models. The lack of support for irregular height grids often results in imported data losing accuracy as it is converted to a regular grid form prior to model construction. Instancing, although supported in VRML, is managed in fundamentally different ways to CAD systems and therefore cannot be exported, often resulting in huge increases in the size of files when exported. Surfaces in VRML are single sided, most CAD systems assume surfaces are double sided. Although most VRML browsers can render surfaces as double sided this results in almost a halving of performance. Concave faces are not yet consistently supported by VRML browsers.

The use of real world co-ordinate systems is also problematic. For example, most VRML browsers are unable to cope with full UK Ordnance Survey co-ordinates due to problems of data length. The internal mathematics used by the browsers is optimised for speed and as a result rounds the user location co-ordinates causing movement to occur in huge steps. Similarly, Z buffer overflow in the rendering engine can also lead to highly unpredictable results. A further limitation is that rotations cannot be used with real-world data as they are based upon rotating relative to the source grid origin, which in most instances is far removed from the area under study. A detailed description of CAD and co-ordinate problems and how they have been worked around is provided in Bourdakos (1996).

As a result it must be concluded that despite being the most commonplace archaeological means for handling and recording spatial information, CAD and GIS system are far from ideal data sources for the construction of VRML models. On the other hand VRML editors themselves fail to provide the

precision and control required for most archaeological modelling tasks. It is still therefore necessary to understand the minutiae of the VRML specifications to be able to both fine tune the source code by hand and to continuously tweak it to ensure compatibility with different browsers.

Two other factors must be considered before investing time and effort in the production of VRML models. VRML is very slow with anything other than the simplest of models, as a result large models are not feasible. VRML is also very difficult to fly. Inconsistent browser interfaces result in a steep learning curve for the pilot, and slow refresh rates can frequently cause disorientation.

Some of these problems will no doubt be solved by the adoption of more powerful machines and the maturing of the software market, others for instance the co-ordinate system problems will not. It must therefore be concluded that VRML cannot be a single answer to all archaeological Virtual Reality requirements, and we must look at other alternative systems.

5 Alternatives to VRML: the role of other systems.

A number of commercially produced Virtual Reality systems and APIs are currently available and as the market matures are being targeted at more specific requirements. For example, Superscape VRT provides a faster operating alternative to VRML for worlds distributed over the Internet, with viewers being freely available, if only on the Intel platform. Accelgraphics Flying Carpet concentrates on pure speed and enables the real time viewing of huge models by incorporating geometric culling and simplification techniques. Such optimisation approaches must be treated with some caution if visibility is to play a significant part in the analysis of the model, but where they are applicable the speed differences can be staggering.

The United States Military Distributed Interactive system (DIS) shares many objectives with archaeology in terms of the representation of complex landscapes, site intervisibility, and the real time control of 'what if' scenarios. Developed at a cost that could never hope to be matched by any non-military research, it must be hoped that now that it is becoming available on a PC platform its potential may be turned towards more positive uses.

If you are modelling in CAD and simply want to look around your model it could be argued that you do not need to translate into any other format at all, thereby avoiding the inherent dangers of data loss or incompatibility. Suitable tools are available to interactively explore CAD models. Micogreen NAVFlyer is one example of a free standing viewer. McNeels Walkabout runs within AutoCAD itself, integrating viewing and modelling to provide the CAD user with a fully interactive development environment.

6 Introducing the 'Bubble World' concept.

In many cases it may be that the freedom of movement provided by traditional Virtual Reality systems is neither necessary nor even desirable. Virtual Reality systems are frequently very difficult to operate, excluding casual or untrained users, and turn all control of the narrative over to the end user. Systems such as apple quick time VR place the user in the centre of a panoramic image wrapped around a cylinder, as do the newly emerging systems which utilise spherical and cubic panoramas, which I will term 'Bubble Worlds'. The latter may serve to provide archaeologists with a critical intermediate alternative between the static image and fully interactive virtual world. In a Bubble World the viewer is free to rotate their viewing angle left/right or/up down, and zoom in or out. Although they do not have unlimited movement within the world they may 'jump' from one pre-set viewpoint to another. Because all of the geometry and rendering is pre-processed a very visually rich model can be generated, including for example ray traced lighting. This can be generated in real time with smooth movement and relatively high resolution on even very low power systems. What is more, the Bubble and Panoramic worlds can be generated in any standard rendering package or directly from photographs, and the two input sources can be easily integrated. If a CAD model has already been built for the purposes of visualisation, a Bubble World can be generated on the basis of it in a matter of hours. One of the first Bubble World systems was VTV developed by Warp. This was used to generate the following examples, all of which were constructed in less than three hours. The reconstructed church (figure 2) was produced from an AutoCAD model by Richard Bayliss (Bayliss forthcoming), the woods at Whittle Dene, Northumberland (figure 3) were generated in VistaPro from a GIS derived data set and the Temple of Mithras (figures 4 and 5) reconstructed from ultra wide angle photography.



Figure 2: A VTV 'Bubble World' based on an AutoCAD model by Richard Bayliss of a Basilica of the Alacami in southern Turkey, rendered in Acurender.



Figure 3: A VTV 'Bubble World' based on GIS derived data of Whittle Dene, Northumberland, rendered in VistaPro.

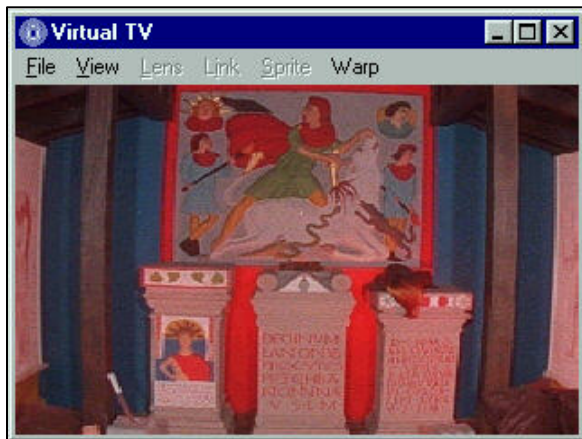


Figure 4: A VTV 'Bubble World' based on a 180° Photograph of the reconstruction of the

Carrawburgh mithraeum at the Museum of Antiquities, University of Newcastle.



Figure 5: A VTV 'Bubble World' based on a 180° Photograph of the Carrawburgh Mithraeum , Carrawburgh, Northumberland.

One very exciting development is the development of RealVR by Live Picture, which extends VRML 2 to enable the mapping of Bubble World backgrounds to VRML scenes and to incorporate high quality image 'sprites' within the scene, in effect combining the speed of Bubble Worlds with the adaptability of VRML. Parallel developments within live pictures will see the incorporation of the "FlashPix" pyramidal image format and "OLIVR" image streaming to permit the delivery of high quality images quickly over the Internet.

This system is causing a lot of excitement and is currently generating more email traffic than VRML despite the smaller installed user base.

7 Conclusions

No single system can meet all of our Virtual Reality needs, Virtual Reality currently pushes desktop computers to the limit of their performance capabilities. If Virtual Reality is to live up to the hype with which it has been presented it is important that it is implemented using the most appropriate system for both the nature of the data and the viewing platform, thereby producing fast usable Virtual models quickly and cheaply.

VRML is important as it represents one of these systems. If you wish to produce Virtual worlds for the widest possible distribution, are prepared to dedicate sufficient resources for development, can live with the inconsistent colour model and do not

need to use inconsistently implemented features no other system can compete. If you do not want to distribute your models over the Internet, have limited

resources, need a consistent delivery platform, high speed or large models, then more appropriate options almost certainly exists.

Bibliography

- Bayliss, R, (forthcoming) The Alacami in Kadirli: Transformations of a Sacred Monument. *Anatolian Studies Vol.XLVII*
- Bell,R, Carey,G and Marrin,C, (eds.) 1996 'The Virtual Reality Modelling Language Specification Version 2.0',The VRML Consortium, URL: <http://vrm1.sgi.com/moving-worlds/index.htm>
- Bourdakis,V,1996 From CAD to VR; Building a VRML model of London's West End. *Proceedings 3rd UK VRSIG Conference*. URL: <http://fos.bath.ac.uk/vas/papers/3rd-UKVRSig/>
- Gillings, M and Goodrick, G T, 1996 'Sensuous and Reflexive GIS: Exploring Visualisation and VRML. *Internet Archaeology 1*, URL: http://intarch.ac.uk/journal/issue1/gillings_index.html
- Gombrich, E H, 1960 *Art and Illusion*. Phalidon
- Mitchell, W L and Pendlebury, M, 1996. Reconstruction of the Egyptian Tomb of Menna using VRML. *Short Paper Proceedings of the 3rd UK Virtual Reality Special Interest Group*. 67-73.
- Nadeau, D, 1997 Visualisation in VRML. URL: <http://www.sdsc.edu/~nadeau/WorldMovers/nadeau.htm>
- Sander, D and Gay, E 1996 VRML based public education -an example and a vision. in VRML Site Magazine, Dec 1996. URL: <http://www.vrmlsite.com/dec96/spot2.html>

Contact details

Glyn Thomas Goodrick
Robinson Library
University of Newcastle
Newcastle NE1 7RU
UK
e-mail: G.T.Goodrick@ncl.ac.uk