

European Research Consortium
for Informatics and Mathematics

ERCIM



Report and Recommendations from the
VRML Workshop 29/30 January 1997,
The Coseners House, Abingdon, UK

Edited by David Duce, Klaus Kansy
and Michael Wilson

Report and Recommendations from the VRML Workshop
29/30 January 1997, The Coseners House, Abingdon, UK

Edited by David Duce, Klaus Kansy and Michael Wilson



1 Introduction

The ERCIM Computer Graphics Network organised a workshop on VRML at the Coseners House, Abingdon, Oxfordshire, on 29/30 January 1997. The workshop was organised by the Steering Committee of the ERCIM Computer Graphics Network in May 1996. It was felt by the Steering Committee that a workshop was needed in order to formulate a European view on VRML. Version 2.0 of VRML was to be published in August 1996. All groups in the ERCIM network showed interest in the topic and a number of key

**Report and Recommendations
from the VRML Workshop
29/30 January 1997,
The Coseners House,
Abingdon, UK**

edited by David Duce,
Klaus Kansy and
Michael Wilson

- 1.1 Cooperative VR
- 1.2 Virtual multimedia presentation
- 1.3 Automatic multimedia generation
- 1.4 VRML for cooperative working
- 1.5 Presentation models and architectures
- 1.6 VR technology and human perception
- 1.7 Semantic preservation in the design of VR presentations
- 1.8 Strengths and limitations of VRML

The key questions to be addressed was "what additions should be in VRML 2.0?"

At that time, VRML 2.0 has been put forward for standardization to the ISO/IEC Committee responsible for computer graphics and image processing (ISO/IEC JTC 1). A review of the document took place in the autumn of 1996 and an early meeting to address the comments raised on the letter ballot took place in the USA in December 1996. Several members of the ERCIM Computer Graphics Network were involved in the preparation of these comments. Three major issues identified in this process were the quality of the VRML document, the suitability of VRML for use in visualization applications and the use of VRML in multi-user worlds.

The meeting opened with an introduction to the background and aims of the workshop, given by David Duce. Anthony Steed from University College London who was a member of the UK delegation to the ISO/IEC VRML Editing

*Report and Recommendations from the VRML Workshop
29/30 January 1997, The Coseners House, Abingdon, UK*

Edited by David Duce, Klaus Kansy and Michael Wilson

11 February 1997

1. Introduction

The ERCIM Computer Graphics Network organised a workshop on VRML at the Coseners House, Abingdon, UK, on 29/30 January 1997. The idea for the workshop came from the Steering Committee of the ERCIM Computer Graphics Network in May 1996. It was felt by the Steering Committee that a workshop was needed in order to formulate a European view on VRML. Version 2.0 of VRML was due to be published in August 1996. All groups in the ERCIM network expressed interest in the topic and a number of topics were identified:

- Cooperative VR
- Intelligent multimedia presentation
- Automatic multimedia generation
- VRML for cooperative working
- Presentation models and architectures
- VR technology and human perception
- Semantic preservation in the design of VR presentations
- Scope and limitations of VRML

One of the key questions to be addressed was "what additions should be in VRML 3.0"?

Since that time, VRML 2.0 has been put forward for standardization to the ISO/IEC Committee responsible for computer graphics and image processing (JTC1/SC24). A review of the document took place in the autumn of 1996 and an Editing Meeting to address the comments raised on the letter ballot took place in the USA in December 1996. Several members of the ERCIM Computer Graphics Network were involved in the preparation of these comments. Three major concerns identified in this process were the quality of the VRML document, the suitability of VRML for use in visualization applications and the use of VRML in multi-user worlds.

The meeting opened with an introduction to the background and aims of the workshop, given by David Duce. Anthony Steed, from University College London who was a member of the UK delegation to the ISO/IEC VRML Editing

Meeting gave an overview of the state of the VRML document within the standardization process. Two talks were given by researchers from the University of Leeds, introducing the topics of VRML in visualization and multi user worlds and describing work in progress at Leeds in these areas. The meeting then split into working groups to address these two topics in more detail. It had been hoped that there would be the opportunity to review the new VRML document arising from the Editing Meeting, but this was not available in time for the workshop. It had also been envisaged that some work could be done on formal description of VRML 2.0, but it did not prove possible to do this and this will remain a task for a future meeting.

At the close of the workshop, reports were received from the two working groups and recommendations to ERCIM were formulated.

A list of attendees at the workshop, together with email addresses is contained in Annex A.

2. The VRML Document

Anthony Steed of University College London, described the current state of the VRML document and the discussion and decisions made at the Editing Meeting in December 1996. Copies of Anthony's slides are appended to the paper version of this report.

Anthony's interest in VRML came from his PhD work, where he developed a data-flow like VR system which had architectural similarities to the sensor - logic - effects model of VRML.

Development of the VRML document appears to have been driven by the annual occurrence of the Siggraph Conference. The VRML Architecture Group (VAG) conducted the discussion of VRML through a mailing list. Two systems were designed to meet the perceived requirements of the community: Moving Worlds from SGI and Active VRML from Microsoft. Sony were involved in the Moving Worlds proposal, but did not make a separate submission. This community process resulted in the choice of the SGI Moving Worlds proposal. VRML 2.0 was published on August 4th 1996 and is available at:

<http://vag.vrml.org/VRML2.0/FINAL>

and extended VRML 1.0 with interaction, animation, scripting and extensibility functionality.

Meanwhile, discussions were taking place between the VAG and ISO/IEC concerning the standardization of VRML by ISO/IEC. Agreement was reached in June 1996 that VRML 2.0 would be brought into ISO/IEC JTC1/SC24, the computer graphics and image processing committee. The VAG would act as architects and the ISO/IEC VRML Rapporteur Group, set up within SC24 to manage the ISO/IEC VRML project, would comment on the document and assist with the editing process.

The document entered the ISO process as a Committee Draft and was circulated for letter ballot.

The UK, along with other nations, submitted detailed comments on the document. The UK disapproved the draft for the following reasons:

- The overall quality of the document, viewed as an International Standard, was poor.
- The document lacked a concepts section.
- There was no explicit specification of the structure of a VRML file, for example, expressed as a BNF specification.
- VRML was unusable across the range of SC24 application areas, for example, there were felt to be deficiencies for handling visualization applications.
- There was no formal process for extending VRML. Statements such as "you can do it through prototypes" or "you can write a JAVA applet to do that" were felt to be too weak in an International Standard.

The outcome of the letter ballot was 6 yes votes, 3 yes votes with comments, and 2 no votes (UK, USA). 500 comments (weighing 3 pounds!) were received with the votes. The Editing Meeting in San Diego in December 1996 was convened to address these comments, produce a Response Document (a formal reply to the submitters of the comments) and instructions for the Document Editors for preparing the next version of the document.

The comments were well-received by the VAG members present at the meeting; an additional 100 comments were added from the VRML "community". There was great reluctance to change any functionality described in the document that was not actually broken. The main reasons given for this were the number of browsers currently available (at least 8) and the number of VRML worlds that have been developed (at least 20,000), not to mention the number of VRML books! However, it was emphasised that the VRML Consortium (more of this later in this report) would encourage extensions to VRML 2.0.

The major changes agreed were:

- Editorial. All comments were processed. The use of language will be made more consistent and a glossary of terms will be added.
- Spelling. Oxford Dictionary spelling will be used except for field and node types (this would affect existing worlds and browsers and was deemed to be undesirable).
- Scope. The Scope clause will be redrafted to set more realistic expectations of where and how VRML is useful. The VRML Consortium are setting up groups to look at extensions to VRML which will aim at extending the scope.
- Concepts. A new concepts clause will be added. Some new architectural figures produced by the UK in their comments will be added (which address the execution model and conceptual processing of VRML data), and the key

concepts time, event processing, lighting, sensors, interaction and navigation, will be introduced. The terms browser, author and user will be defined.

- Interaction. Modifications will be made in line with the UK comments.
- Grammar. An improved grammar has been drafted. A BNF definition was seen as desirable, but concerns were raised about ensuring that the grammar is accurate and faithful to the VRML description.
- Conformance. The base requirements for VRML (essentially what a browser has to guarantee to parse) will be collected in one place. File and capability limits will be added, effectively these will be the lowest common denominator of the limits and capabilities of current browsers.

Other comments did not receive such a sympathetic hearing.

- The request for a separation between the clear text encoding and the abstract VRML syntax specification was rejected.
- The request for richer data types was rejected, though ranges will be added to data types.
- The need for cross-platform support for VRML (which includes the strong requirement that it should run on a PC) means that colour, fog, lighting models, etc. have to be kept simple, otherwise browsers running on small machines would have no hope of rendering VRML worlds.
- Absolute reference frames for time and distance are needed in VRML to ensure composability.
- The observation was made that all the extensions proposed in the comments "could be prototyped".

Anthony's talk gave rise to much discussion. The key points raised are summarized below.

1. Scripting. There are some signs of convergence between VRML Script and JAVA Script. There are no plans in this phase to standardize an external authoring interface; it is possible that Sony and SGI will rewrite and resubmit their proposals. The Sony view of scripting seems to be very different to the SGI view. A browser can be VRML compliant without supporting scripting. Concerns were raised that recourse to scripting is necessary in order to do some apparently rather elementary things; this seems to point to some deeper deficiencies in the standard.
2. World View and Liquid Reality browsers support JAVA. Most of the non-SGI browsers look as if they will support JAVA. SGI provide a scripting interface.
3. Routes to modification. Two routes appear to be open:
 - New work within ISO/IEC
 - New work through VRML Consortium

The VAG has now dissolved itself; a new VRML Consortium (modelled on the X Consortium) has been formed. This Consortium has members (currently around 30, cost is \$1,000 for non-voting, \$7,500 for voting and \$15,000 for chartered members). It is believed that individual members (non-voting) might also be admitted. The Consortium has Working Groups, which

it appears anyone can join. Some organizations represented around the table were considering joining the Consortium.

4. Microsoft have apparently announced that they will be making the source code of World View and Liquid Reality browsers available as reference implementations.
5. The VRML document will contain informative annexes on scripting and JAVA interfaces. The issue of the extent to which scripting/JAVA are necessary was a recurring topic during the workshop.
6. A binary encoding of VRML is very likely to materialize, but a clear text encoding is not likely.
7. It was noted that in many ways the best way to influence the VRML community is by developing and making available, "working code".

3. Multi User VRML

Stuart Lovegrove, a PhD student from the University of Leeds described the research he is planning to do. His thesis title is "Sustainable Communities for Collaborative Research" and he started work on this topic in October 1996, so the work presented is very much preliminary work and first ideas.

The background to Stuart's work comes from the COVISA project at the University of Leeds School of Computer Studies. The COVISA (Co-Operative working in Visualization and Scientific Analysis) project looked at understanding requirements for group working in visualization. Fully working collaborative extensions of the IRIS Explorer modular visualization system have been produced. These allow several users to collaborate by exchanging data or by joint control of the user interface.

The development of WWW technologies have also lead the group to look at using IRIS Explorer as a visualization server. The user interface is a HTML page, and the server returns the visualization generated as a VRML file. Collaborative working is achieved by storage of previous results. Further details of this work were given in Jason Wood's talk, see section 5.

The aim of Stuart's project is to combine these results, to provide WWW visualization and multiuser, interactive and collaborative working in real time, along with persistent data storage. The initial architecture is based on a central server connecting to user and data clients. User clients receive VRML files and JAVA. The central server controls message passing between users. It is being written in JAVA (to aid portability) and supports multithreading (allowing connection requests to be serviced at any time) and broadcasting (allowing incoming messages to be sent to all connected users).

VRML files are used to represent shared scenes and avatars. Currently these are statically defined. The VRML structure incorporates a script node with an external JAVA class. VRML routes are run from the script file to other nodes within the shared scene. Currently routes are also statically defined. Eventually routing needs to be tailored to what each particular user can interact with.

There is a socket interface between the JAVA client and the network, through which events are passed from the server to the VRML file and the VRML file to the server. The data client also has a socket interface to the server. Users moving or interacting in a scene send messages to the server, which provides a method of viewing all events.

Some proof of concept experiments have been carried out, based on the Dimension X Liquid Reality browser. Two simple scenarios were shown. The first provides a shared pointer, multiple avatars representing users and a "clickable" cube which changes its colour in response to clicks. The demonstration correctly displays the views of user 1 looking at user 2 and user 2 looking at user 1 on their respective browsers. In the second scenario, users are placed in a simple world. Avatars now have a humanoid appearance.

Stuart stressed that the work reported represents the very first stage in his project and there are many questions and problems still to be addressed. Issues raised include:

- Portability
- Efficiency/speed
- Consistency
- Persistence
- Technical capability
- VRML capabilities
- Creation
- Scope
- Services
- Compatibility

4. Examples of the Usage of VRML and VRML Worlds

Michael Wilson gave a short presentation on current uses of VRML in a range of different application domains and examples of applications that VRML ought be able to support.

The Community Place system uses an architecture similar to Stuart's, except that the VRML and database servers are combined. Position/location is broadcast from avatars; avatar motion can be expressed.

The Sony circus park test demo is an illustration of the games sector.
(<http://vs.sony.co.jp/index-e.htm>)

The DIVE 3D VR system does not use VRML but is an example of a multi user VR application which VRML should support.
(<http://www.sics.se/dce/dive/>)

Work at RAL on the use of VRML to visualize engineering design is an illustration from the engineering sector. The RAL group have generated VRML from a CAD package, in order to explore the assembly sequence of a complex object (a very large detector for a high energy physics experiment at CERN). (A fuller description will appear under <http://www.ccl.ac.uk>)

A Cambridge company have developed a web agent, Autonomy, which uses 2D icons to visually represent agent state, again this is an example of an application which VRML should support. (<http://www.agentware.com/>)

The QPIT project at Lancaster University has developed an approach to multi-user embodiment within data visualizations - populated information terrains. This again is an example of a type of application that VRML should be able to support. (<http://www.comp.lancs.ac.uk/computing/research/cseg/projects/qpit>)

The Persona project at Microsoft is looking at general Lifelike Computer Characters. (<http://www.research.microsoft.com/research/ui/persona/related.htm>)

5. Visualization over the World Wide Web - A Server-Side Approach

Jason Wood gave a presentation about work carried out by himself, Ken Brodlie and Helen Wright at the University of Leeds.

Jason described a VRML 1.0 system for server-side visualization across the web.

A traditional model of visualization considers the transformation of data to image in three stages: filter, map and render. Overlaying this with a server-client relationship, the server may be thought of as a publisher and the client as a user.

A client-side approach associates the data with a WWW server and the transformation process (filter, map, render) with the client. Data are requested from the server (e.g. through a form interface) and are returned for local processing by the visualization system located with the client.

A server-side approach associates the data, filter and mapping transformations with the publisher in a visualization server, and the rendering process with the client. This gives control for data mapping to the server rather than the client. Again a forms interface to the visualization server can be used to define the type of transformed data required by the client.

Leeds have developed such a visualization server, based on IRIS Explorer. A forms interface is used to describe the visualization required. This is passed through a CGI script to the visualization server which generates a script to drive IRIS Explorer to produce the visualization required. The geometry file output by IRIS Explorer is captured and translated to VRML and passed back to the client.

An example application was described. Air quality data are available from a public source. Data are collected hourly and made available about 2 hours later. Currently the data supplier provides the latest hourly data values, previous data values, some forecasts of future values and some basic, predefined graphs. A web-based interface has been written to this data service, using the Leeds visualization server. From the interface form, the user can select the site, pollutant, time etc. for which data are required. The server returns a VRML file containing a visualization of the selected data.

The work has revealed some limitations of VRML 2.0, for example: adding a colour scale so that the user can always see it is not at all easy. There is a solution using proximity and sensors, but this seems an overly complex mechanism for satisfying a rather simple requirement. Solutions such as putting the colour scale in a separate viewing area do not work because one cannot guarantee consistency of colours over areas.

6. Working Group Reports

6.1 Multiuser Virtual Worlds (MUVW)

This group was chaired by Michael Wilson, CCLRC.

The group viewed the basis of discussion as being four document sets :

1. The body of the VRML 2.0 specification which was taken as finalised as an ISO standard to which all VRML 2.0 tools must conform.
2. The appendixes to the standard on JavaScript, Java and external authoring interfaces to VRML2.0 which is only informative to those developing VRML2.0 compliant tools. The consequence is that some tools adopt one or other interface, so that code written using only one interface will not be usable in browsers or tools which do not support that interface. Sony and Microsoft appear to be supporting the Java interface, while SGI are supporting the JavaScript interface. Even then Sony and Microsoft browsers use different Java engines and may not interpret the Java in the same way. SGI have browsers for the SGI and PC platforms which do not operate in the same way either - this issue is returned to below under Portability.
3. The proposals for multi user virtual worlds before the VRML Consortium working group on MUVW as layers above the existing VRML 2.0 standard for developing MUVW. These proposals accept the standard as is, and do not modify it; they merely add layers of standardisation on top of it for MUVW.
4. The proposal for Universal Avatars currently before the VRML Consortium Avatar Working Group. This also takes the VRML 2 standard as is, and proposes a layer on top of it. This proposal was not discussed in detail, although it was circulated at the working group to participants for consideration.

The working group was very much informed by the view of the current state of VRML and the MUVW proposals of Anthony Steed (UCL) and Stuart Lovegrove (U of Leeds) whom the group thank for their contribution.

Scope

It was agreed that the scope of MUVW would be expected to include many things (everything that was suggested) :

1. Multi user meetings/office/CSCW environments - meeting room environment supporting user avatars, whiteboard, digital video videoconferencing interfaces, pointers and other object services.
2. Envisionment of faces for multi user teleconferencing - combining mouth and eyes from MPEG with animation for the rest of the envisionment.
3. Multi-User Games - this is clearly the target market for Sony and others in the development of VRML 2.0 so we expect these needs for presentation, interaction and payment to be met.
4. Real Work Tasks - For example the planning of spacial activities in the VR before their enactment in the real world; medical surgery, engineering construction etc ...
5. Education - Multi user seminar, tutorials simulations etc for educational exploration and interaction.
6. Data-spaces such as the WWW or databases can be spatially visualised and then users envisioned at locations in them to show the information they are/have acquired. Multi-user support here could lead to chat meetings to elucidate the visualisation between those experienced in the same information.
7. Intelligent agents as well as human avatar envisionment in all of these spaces - human envisionment requires control of the human avatar through input devices, agent envisionment requires control through large software programs including generation of sound, action which would be caught by the human input devices and transmitted for human user embodiment.

Suitability of VRML2 for MUVW

The experienced view was that DIVE (from SICS) was the preferred MUVW tool today rather the VRML2 ones, but it is supported on UNIX and not PC platforms since it passes massive object state changes between peer to peer connected clients, and includes no security so one user can destroy another's VW. Therefore it is acceptable for responsible research use (it is also free with source code), but impractical in many ways for commercial use.

VRML2 was accepted as the future for MUVW and there was not seen to be any problems in it for most uses of rendition - most of the discussion was about the interfaces, tools, and architectures for using VRML as the graphics language about documents 2 and 3 mentioned above.

Recommendation :

The changes requested to the VRML standard itself were :

1. The first change requested is for Intelligent Agents to produce sound output dynamically which is not supported by the current standard since they have to produce WAV files which are then passed to the VRML2 tool; this works but is a clumsy hack.
2. The second change requested was for the creation and deletion of objects dynamically in the world. Creation can be currently done by creating strings and then turning these into objects, which is a real hack; deletion is not entirely possible because of the need to maintain different users worlds consistently, so some remnant of the object exists even if one user's Java engine has garbage collected most of the object. Deletion problems make it impossible to leave a MUVW meeting since the avatar cannot be deleted.
3. The third change requested was to register object/agent to object/agent collisions since only object to environment collisions are currently registered. In a MUVW with many active agents or objects this is a requirement that seems strange not to have been met in the current standard.
4. The fourth change was to introduce ever present menus/controls into virtual worlds. These can be emulated through sensor driven billboards which are always turned to the user, but these must be moved through the world to follow the user's movement in order to always stay visible. An explicit menu node which performed this functionality would make life easier.

Architectures for MUVW

Suggested architectures include the client browser to HTTP server link with the downloading of a MUVW, then the linking to the MUVW server to support the existence, location, and action of objects in the MUVW such as other users. Clearly the main role of the architecture is to maintain consistency between different user's views on the MUVW.

1. The first main distinction in architecture is whether to use peer to peer multicasting (e.g. SICS DIVE, and Superscape) of state between clients or to maintain the single server (e.g. Sony Community Place) which broadcasts to all clients. The server communication will become the bottleneck in the server model. MUVW examples from the game Quake and Superscape suggest limits of 16 and 25 users respectively in a MUVW before 28.8Kb/sec modem communications (the target platform for VRML usage) make the MUVW unusable due to the server communications bottleneck. In contrast, charging models are much easier to place on server based systems than peer to peer ones. Sony appear to be suggesting a solution which is to introduce Zones in the MUVW as part of the Living Worlds proposal so that a server broadcasts to clients in a zone, and multicasts moves between zones to other servers. This limits the number of users in a zone, but a zone could be a small space in the MUVW. A third variant of this is to allow each client access to a local database, and allow the databases independent control of how they maintain consistency between them.
2. The second major architectural distinction is whether to place functionality in the browser or in the MUVW. If the functionality is placed in the VW, then it is portable but larger and hence slower to load and since it is interpreted JAVA, slower to execute than if placed in the browser as compiled

C++. Avatar controls could be placed in the MUVW and presented as sensor driven billboards (a hack to emulate menus) to allow special avatar actions to exist in some worlds. The tradeoff is between portability and performance; Sony have placed the MUVW support for avatar control and MU-Chat (voice and text) in the browser with the resulting lack of portability of their worlds to other browsers - this issue is returned to below under Portability.

3. The third distinction is in the messages that are passed themselves in order to maintain consistency. Four options were discussed :
 - sensor events at client are transmitted to server and broadcast with the result that different browser implementations produce different world states due to service interaction. This appears to be ruled out by the Living Worlds proposal.
 - Pass the visible object state to maintain object consistency; this is used in DIVE and Superscape through peer to peer multicast communication.
 - Pass a semantic message - appears to be the most common procedure, but the set of messages is not standardised at all.
 - Rely on a database at each client browser to maintain consistency with its peers.

The Living Worlds proposal does not address these three possible alternatives.

Efficiency/ Technical Capability

The efficiency of the architecture will be limited by a bottleneck in any of 3 places:

1. Network Communications. VRML2 is aiming to support MUVW through 28.8 Kb/sec modems to end users. The more, larger packets used in the message passing, the more strain is placed on the network. Server end communications appears more likely to be the rate limiting factor than client end. VRML2 is a bulky ASCII format, and a more compact binary format would place less strain on the communications. Gzip and Gunzip at each end of a world download adds to the time at each end although resulting in a transmission saving. This is not enough and a binary format is required as one of the VRML consortium working groups is currently developing. The extreme view on this position has been attributed to Jon Crowcroft at UCL who would like to treat VRML as a datatype, and introduce socket level rules to handle this specific datatype for packet loss, packet re-ordering etc as a result of network performance (Quality of Service).
2. Browser. VRML2 is aiming to support MUVW in desktop PCs for end users. Introducing Java or JavaScript API introduces an overhead at interpreting packets over lower level protocols. Generation of sound, intelligent agent planning could grind the client machine if these are included.
3. Server. If the server processes multiple services (e.g. collision, object services as described in the consistency section below, and see also the services section below) then it could become a processing bottleneck. The simpler services on the server, the less processing. There is no guidance on placement of services on client/server in the Living Worlds proposal.

Recommendation :

The MUVW proposal (documents in group 3) should provide a server API to register services, and guidance on the placement of services at client or server.

Object Creation by Users

VERMELgen uses Jverge Java class libraries to support the construction of a V W from inside another VW, but uses the stand alone Jverge modeller which may not be strictly VRML2 compliant.

There are no proposed mechanisms for MUVW users to create or import previously created objects into MUVWs. This is a problem where a user at a meeting wishes to introduce a model to explain something.

This issue is confounded by the general problems associated with object creation in VRML2 itself.

Recommendation :

The MUVW proposal (documents in group 3) should address this issue and propose standards with interfaces.

Consistency

Maintaining consistency between different user views of the MUVW was addressed above in architectures. The main problem that most approaches have is that of apparent simultaneous actions. For example, if two users simultaneously pick up a wallet, which one has it. The window of a simultaneous event is controlled by the delays in the network communication. These are minimised if each object has an object server thread running on some machine in the MUVW, and this thread is accessed before the world for any user is updated (therefore two users can never pick up the wallet). However, the cost of this solution is an increase in processing, and another communication overhead. In a client/ single MUVW server architecture this is putting more strain on the server communications bottleneck; in a peer to peer architecture the threads for different objects can be spread throughout the participating peers to reduce the bottleneck problem, but the communication still includes overheads.

Recommendation :

The MUVW proposal (documents in group 3) should address this issue and propose standards with interfaces.

Persistence

World Persistence : A user leaves a MUVW, but other users remain in it, and make changes to it. A user then joins the world: do they enter at the normal

starting state, the state at which they left it, or the state in which other users see it? Presumably, as a user enters a world and registers with the world server, the server knows the start up state, and the current state of the world, and informs the user's browser of the difference to update the user's view of the world to the current state and ignores the state at which the user left it. No document addressed this issue.

User Entry Point Persistence : In VRML2 worlds users enter at a standard entry point. If the user has left a world he will not re-enter at the exit point, but always at the standard entry point. If the local browser stores the exit point from each world a user has left, then re-entry at that point could be supported.

Recommendation :

The MUVW proposal (documents in group 3) should address this issue and propose standards with interfaces.

Security & User Authentication

It is necessary to limit areas of MUVW to users, costing models require that all users are identified, and it is required that users cannot destroy other users' worlds.

Living Worlds proposes zones where scripts can be attached at the boundary to enforce security checks on entrance, but no procedures for this are included at present. No general user identification model is proposed. Object destruction is not supported in VRML, so it is not a problem here.

Recommendation :

The MUVW proposal (documents in group 3) should address this issue and propose standards with interfaces.

Device Drivers and Rendering Hardware

All current browsers support non-immersive MUVW assuming desktop use. To extend a browser to a new display device, to add new input devices (e.g. gloves, game controls etc.) or specifying the parameters of the output device in order to tune the presentation requires getting inside the browser. No browser currently publishes its API to support this. The Liquid Reality browser is issued as Java bytecode so it is possible to decompile this, edit it and then recompile it, but this would be a breach of the license and is in any case hardly an answer to the problem.

Recommendation :

Browsers publish API's to support the change in I/O device and properties.

Services

Services such as :

Avatar, Pointers, Whiteboard, link-list, DB I/F, Search, Sort, Security

could be provided by server or browser. If they are to be supported by browsers to improve efficiency then we need to be able to add to them rather than limiting them to browser supplied services. For both server and browser then we need the API to :

1. Register services with name, class, language (e.g. Java or Javascript)
2. List registered Services, providing calls
3. Invoke Service.
4. Close Service.
5. Remove Service.

Recommendation :

The MUVW proposal (documents in group 3) should address this issue and propose standards with interfaces. Browsers and Servers should publish API's to support services.

Avatars

The avatar proposal appears to support one avatar per browser in order to overcome portability problems - this is not the general solution to this problem.

It should be possible to present different avatars in different worlds (e.g. one for flight simulation, another for business meetings, etc.) which does not appear possible in the proposal. There are potential problems associated with introducing multiple avatars of the same person into a single world and managing the control of these.

There is no guidance in the Universal Avatar proposal as to embodiment as avatars for different tasks.

Recommendation :

The Universal Avatar proposal (documents in group 4) should address this issue and propose standards with interfaces.

Portability

The major problem with using VRML2 MUVWs today is the lack of consistency in the browsers. MUVWs only run in the browser for which they were written. VRML1 worlds run in all VRML1 browsers, so the worlds are portable, but until there is more portability resulting from standardisation, VRML2 MUVWs are

too much tied to the browser, so we expect a browser war until somebody produces a browser which is available on all platforms that can break the market.

The problems of Java versus JavaScript versus Authoring Interface conformance to different normative appendices were mentioned above. Similarly, the issues of placing functionality in the browser versus the MUVW itself. Different server structures are mentioned under the architecture heading.

Recommendation :

Conformance to standards for VRML2 browsers, servers, MUVW protocols is required to enable portability of MUVWs.

6.2 Visualization Working Group

This group was chaired by Klaus Kansy, GMD.

Is VRML a Visualization Tool?

The authors of VRML claim that it can be used for visualization. At the workshop, examples were seen from the University of Leeds and British Telecom which use VRML in visualization applications. VRML was seen to be useful for some kinds of visualization, but it does not address or solve many problems.

Advanced visualization needs modelling power. It is not sufficient to transfer low level data and simple scripts but compact descriptions at a high conceptual level and powerful operations at that level are required. The description of the data to be visualized may take the form of a set of values over a regular or irregular grid, or may be expressed analytically, for example, by a set of differential equations. Descriptions at this level are beyond the scope of VRML. Throughout the meeting there was the constant reminder that one of the design goals for VRML was that it must run on a PC.

Visualizations often involve large volumes of data. Medical imaging applications, for example, frequently involve image sizes in excess of 10 Mbytes. There is a need for some form of lazy down-loading in VRML, so that data are only requested when they are needed, for example, depending on the user's viewing direction. This is a non-trivial problem, which is not adequately addressed by level of detail. There is a need for both server and browser intelligence in order to decide what needs to be down-loaded and when.

A Reference Model for Visualization

In his introductory talk, Jason Wood, University of Leeds, had introduced the following reference model for the visualization pipeline:

data ----- filtering ----- mapping ----- rendering ----- image

This visualization reference model was discussed in detail. The filter process selects and transforms application data. The map process maps the data to be visualized onto visualizable elements and render generates the image to be viewed. Although there were discussions about the limitations of this model, it was felt overall that it is helpful and assisted in structuring some of the discussion.

Where is VRML located in the pipeline in this reference model? It was felt that VRML belongs between the mapping and rendering processes. The VRML representation contains the data the browser needs for rendering this VRML world. Views of the visualization from different view points can be computed by the renderer, though in some circumstances this may require regeneration of the model by the application. Consider a mesh containing pressure data at points on the mesh. To compute the pressure at intermediate points requires interpolation of the data. Suppose pressure values are mapped on to colour values. A renderer displaying this image without reference to the data will compute colour values at intermediate points on the grid by interpolating in colour space, rather than interpolating in data space and performing the data to colour mapping. Such different interpolation schemes can generate very different results.

Limitations of VRML

It is difficult to build menus into a VRML scene. Are they part of the 3D world? Adding annotations or colour scales is also difficult. VRML does not currently allow multiple views in the same window. Adding this functionality would start to address these concerns.

Solid Modelling

The group also looked at the needs of solid modelling, which turn out to be somewhat similar to the needs of visualization. Model description, interaction and lighting are key concerns. As in the case of visualization, there are examples where it is necessary to go back up the pipeline from the viewer to the publisher in order to carry out certain operations.

Solid Modelling

Scripting was hotly discussed. The by now standard answer "scripting solves it" was seen to be very unsatisfactory. The problem is that scripting is very important and hence ought to be a part of VRML, yet support for scripting by browsers is optional. Scripting ought to have a more official status.

Conclusions

Some conclusions were reached about missing functionality needed to make VRML a useful tool for visualization tasks.

1. Multiple views in the same window should be added.

2. Lazy down-loading is necessary for many applications with high data volumes.
3. In order to control in which space interpolation happens, it is necessary for the viewer to be able to map application data on to colour values. Addition of a primitive such as "indexed face set with data" would address a wide range of requirements.
4. A more appropriate colour model is needed to meet the requirements of colour handling in visualization, e.g. the absolute (device independent) specification of colour.
5. Sound is mentioned in VRML but acoustic properties are not. Diffusion, refraction and transmission acoustic properties should be added to the material properties in VRML, to enable realistic 3D sound rendering.
6. Sound generation can only be done from a file in VRML, analogously to texture in graphics. "Sonification" (the acoustic analogue of visualization where sound is generated from an abstract description or even arbitrary data) is a field of growing importance. There is a need to add interfaces to sound synthesis.
7. Textures are supported but pixel and voxel-based geometry nodes are missing which are important methods for visualization, for example in medical imaging.


7. Recommendations

The closing plenary session, after receiving reports from the Working Groups discussed the recommendations that should be made to ERCIM, through the ERCIM Executive Committee. VRML is a development of importance to many of the ERCIM institutes and a development that ERCIM should seek to influence through its own offices as well as through those of its own member institutes.

1. ERCIM as an organization should join the VRML Consortium. The fee for a voting member is \$7,500. The VRML Consortium is at an early stage and it would make sense for ERCIM to join as a single entity rather than for each member to join individually.
2. Developments of VRML are formulated and discussed within Working Groups in the VRML Consortium. In the multiuser extensions area there are existing Working Groups that are addressing some of the issues raised at this workshop. Reports of the Multiuser Working Group discussion should be fed into these groups as appropriate.
3. There is no existing Working Group in the VRML Consortium in the visualization area. ERCIM, having joined the VRML Consortium, should propose the creation of a Visualization Working Group. A working group needs to identify a chairman and secretary/webmaster. Julian Gallop, RAL, expressed interest in taking on the Chairmanship role. The ERCIM Computer Graphics Network will supply a person for the secretary/webmaster role.

Annex 1 - Lst of Attendees

Boyd	David	RAL, UK d.r.s.boyd@rl.ac.uk
Clark	Adrian	University of Essex, UK alien@essex.ac.uk
Collins	Lewis	British Telecom, UK lewis.collins@bt-sys.bt.co.uk
Crossley	Martin	British Telecom, UK martin.crossley@bt-sys.bt.co.uk
Drumm	Ian	University of Manchester, UK zzcguid@afs.mcc.ac.uk
Duce	David	RAL, UK dad@inf.rl.ac.uk
Ducrot	André	INRIA, France andre.ducrot@inria.fr
Francis	Alan	Page Description Ltd, UK a.h.francis@open.uk
Gallop	Julian	RAL, UK jrg@inf.rl.ac.uk
Herman	Ivan	CWI, Netherlands ivan.herman@cwil.nl
Hopgood	Bob	RAL, UK frah@inf.rl.ac.uk
Jeffrey	Andy	University of Manchester, UK andy.jeffrey@mcc.ac.uk
Joan-Arinyo	Robert	UPC, Spain robert@turing.upc.es
Kansy	Klaus	GMD, Germany kansy@gmd.de
Kennaway	Richard	University of East Anglia, UK jrk@sys.uea.ac.uk
Lovegrove	Stuart	University of Leeds, UK stu@scs.leeds.ac.uk
Martinho	Carlos	INESC, Portugal cma@minerva.inesc.pt
Mueller	Markus	AED, Germany mueller@aed-graphics.de
Rist	Thomas	DFKI, Germany rist@dfki.uni-sb.de
Sastry	Lakshmi	RAL, UK lakshmi.sastry@rl.ac.uk
Solano	Lluis	UPC, Spain solano@turing.upc.es
Steed	Anthony	University College London, UK a.steed@cs.ucl.ac.uk
Wilson	Michael	RAL, UK mdw@inf.rl.ac.uk
Wood	Jason	University of Leeds, UK jason@scs.leeds.ac.uk



The European Research Consortium for Informatics and Mathematics (ERCIM) is an organisation dedicated to the advancement of European research and development, in the areas of information technology and applied mathematics. Through the definition of common scientific goals and strategies, its national member institutions aim to foster collaborative work within the European research community and to increase co-operation with European industry. To further these objectives, ERCIM organises joint technical Workshops and Advanced Courses, sponsors a Fellowship Programme for talented young researchers, undertakes joint strategic projects, and publishes workshop, research and strategic reports as well as a newsletter.

ERCIM presently consists of thirteen research organisations from as many countries:



**Central Laboratory
of the Research
Councils**

Rutherford Appleton
Laboratory
Chilton, Didcot
GB-Oxon OX11 0QX

Tel: +44 1235 82 1900
Fax: +44 1235 44 5385
<http://www.cis.rl.ac.uk/>



**Centrum voor
Wiskunde
en Informatica**

Kruislaan 413
NL-1098 SJ
Amsterdam

Tel: +31 20 592 9333
Fax: +31 20 592 4199
<http://www.cwi.nl/>



**Consiglio Nazionale
delle Ricerche**

IEI-CNR
Via S. Maria, 46
I-56126 Pisa

Tel: +39 50 593 433
Fax: +39 50 554 342
<http://bibarea.area.pi.cnr.it/ERCIM/welcome.html>



**Czech Research
Consortium
for Informatics
and Mathematics**

FI MU
Botanicka 68a
602 00 Brno

Tel: +34 3 401 72 89
Fax: +34 3 401 62 10
<http://www.utia.cas.cz/CRCIM/home.html>

DANIT

**Danish Consortium
for Information
Technology**

CIT
Forskerparken
Gustav Wieds Vej 10
8000 Århus C

Tel: +45 8942 2440
Fax: +45 8942 2443



FORTH

**Foundation
of Research
and Technology -
Hellas**

Institute of Computer
Science
P.O. Box 1385
GR-71110 Heraklion,
Crete

Tel: +30 81 39 16 00
Fax: +30 81 39 16 01
<http://www.ics.forth.gr/>



**GMD -
Forschungszentrum
Informationstechnik
GmbH**

Schloß Birlinghoven
D-53754 Sankt
Augustin

Tel: +49 2241 14 0
Fax: +49 2241 14 2889
<http://www.gmd.de/>



**Institut National
de Recherche
en Informatique
et en Automatique**

B.P. 105
F-78153 Le Chesnay

Tel: +33 1 39 63 5511
Fax: +33 1 39 63 5330
<http://www.inria.fr/>



**Instituto
de Engenharia
de Sistemas
e Computadores**

Rua Alves Redol 9
Apartado 13069
P-1000 Lisboa

Tel: +351 1 310 00 00
Fax: +351 1 52 58 43
<http://www.inesc.pt/>



**Swedish Institute
of Computer Science**

Box 1263
S-164 28 Kista

Tel: +46 8 752 1500
Fax: +46 8 751 7230
<http://www.sics.se/>



**Schweizerische
Gesellschaft
zur Förderung
der Informatik und
ihrer Anwendungen**

Math. Dept.,
ETH-Zentrum
CH-8092 Zürich

Tel: +41 1 632 22 25
Fax: +41 1 632 10 85
<http://www-dbs.inf.ethz.ch/sgfi/>



**Stiftelsen
for Industriell og
Teknisk Forskning
ved Norges Tekniske
Høgskole**

SINTEF Telecom &
Informatics
N-7034 Trondheim

Tel: +47 73 59 30 00
Fax: +47 73 59 43 02
<http://www.informat-ics.sintef.no/>



**Magyar Tudományos
Akadémia -
Számítástechnikai és
Automatizálási
Kutató Intézete**

P.O. Box 63
H-1518 Budapest

Tel: +361 166 5644
Fax: +361 166 7503
<http://www.sztaki.hu/>



**Technical Research
Centre of Finland**

VTT Information
Technology
P.O. Box 1200
FIN-02044 VTT

Tel: +358 9 456 6041
Fax: +358 9 456 6027
<http://www.vtt.fi/>

**Central Office:
ERCIM**

Domaine de Voluceau, Rocquencourt, B.P. 105, F-78153 Le Chesnay Cedex, FRANCE

Tel: +33 1 39 63 53 03 Fax: +33 1 39 63 58 88

<http://www-ercim.inria.fr/>