Implementing Multi-user Virtual Worlds: Ideologies and Issues

Chris Greenhalgh,[1]
University of Nottingham, UK

Abstract

This paper is a personal reflection on the different communities that are contributing to the design and implementation of multi-user virtual world systems. Given the Web3D Consortium's considered re-labeling from "VRML" to "Web3D" this seems an ideal opportunity to take stock of some of the interests which have historically fallen at the edge of (or outside the scope of) specifically VRML activity. I discuss a number of distinct communities which have been developing characteristic multi-user 3D technologies: Internet virtual worlds (including VRML); distributed simulation; games; visualization; and "pure" multi-user VR research. In each case I will suggest the primary concerns and goals of each, as well as sketching out a little of their respective backgrounds. I then suggest a taxonomy of approaches to multi-user virtual reality, which can inform reflection and future work, and consider the extent to which these communities may be converging. Finally I make some suggestions for next steps in VRML multi-user support. It is my hope that this paper will give the reader an entry point into some of these related communities, and will facilitate the establishment of new synergies within the community at large.

CR Categories and Subject Descriptors: I.3.8 [Computer Graphics]: Applications; I.3.6 [Computer Graphics]: Methodology and Techniques – Interaction Techniques; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism – Virtual Reality.

Additional Keywords: Multi-User, Networked Apps.

1 INTRODUCTION

Multi-user 3D virtual worlds have the potential to address a very broad range of applications, including entertainment, simulation, visualization, and tele-communication. They are also technically challenging and a productive research area for many disciplines, including distributed systems, human-computer interaction, computer-supported cooperative work, and social studies. Each application domain and research area tends to support its own community of visionaries, researchers and users (actual or projected), and it is often difficult to establish mutual understanding, appreciation and cooperation between these communities.

Given the VRML Consortium's self-conscious re-invention as the Web3D Consortium, this seems to be a good time to invest some time and energy into appreciating and understanding the different communities and ideologies present and active in the area of multi-user 3D virtual worlds. This paper is my own contribution to this effort. In part it is a reflection on the nature of and reasons for our own historically distance from the VRML effort. It is necessarily over-simplified and subjective; apologies to any who feel misrepresented or unjustifiably pigeon-holed. None the less, I think it gives a useful glimpse of the larger picture of evolving activity, and may suggest particular areas of synergy or cooperation which should be taken forward.

The following sections consider in turn five approaches to multi-user 3D technologies, which I have labeled as: Internet virtual worlds (including VRML); distributed simulation; games; visualization; and "pure" multi-user VR research (my own community, and one which I have trouble naming). In each case I will suggest the primary concerns and goals of each, as well as sketching out a little of their respective backgrounds. The penultimate section of this paper presents a taxonomy of approaches to multi-user virtual reality, which can inform reflection and future work, before considering the extent to which these communities may be converging. One might identify a further community primarily concerned with networking in itself, using multi-user VR a driving application; this group would span all of those considered and will not be considered individually in this paper.

2 INTERNET VIRTUAL WORLDS

I take this community to be those individuals, groups, and organizations whose primary concern is to provide access to multi-user virtual worlds for today's "normal" Internet users: typically accessing the Internet from a commodity PC over a dial-up modem connection at a few tens of kilobits per second. I consider the majority of the historical VRML effort to have been focused in this domain. This is certainly the operational domain of the current VRML-based (and VRML-like) multi-user offerings, such as Blaxxun [2] and VRTelecom [3], and the almost stalled Living Worlds Web3D Working Group.

The main goal of this group, in my simplified consideration, is for people – ordinary, "real", people – to use the systems; the more people, the better. Ideally they would also pay for the pleasure, but that is turning out to be easier said than done.

The ideology (axioms?) of this group has the following key components (you are welcome to disagree, of course):

- It must work right now, on the machines, networks and (sometimes) software that people already have.
- Development must be incremental and evolutionary.
- Reliability is important (at least in principle).
- It must look attractive, and be easy to use.

[1]School of Computer Science and IT
Email: cmg@cs.nott.ac.uk

- It must be readily authored.
- It would really help if it were standardized.

These factors, especially the first two, have tended to give rise to an approach or style which starts with single user 3D graphics (e.g. a single user VRML browser) and introduces elements of networked functionality, such as avatars to represent other users, or custom-coded network interactions (e.g. distributed triggering of animations). This as-required approach has tended to favor simple distributed approaches, e.g. using a single central server and TCP (unicast) communication.

So, to a first approximation, good points about the Internet virtual worlds community's work are that there is a great deal of experience with "real" use, a customer-facing orientation, and good links with the art and design communities. On the down side, there is a tendency to premature standardization (*de jure* or *de facto*), and a reluctance or inability to innovate or experiment on a significant scale.

There are some notable exceptions to my generalizations in this category, which should perhaps be viewed as spanning both this and the research category; these include Sony's Virtual Community work, and MERL's Open Community work. I will return to these when I deal with the "pure" research community.

3 DISTRIBUTED SIMULATION

There is a long history of constructing and using multi-user virtual worlds in the (primarily military) distributed simulation community. This community is focussed on the use of multi-user virtual worlds to create realistic and valid training exercises (e.g. ground and/or air and/or sea combat), where real-world training is not possible or appropriate (e.g. too expensive, too messy, too public, or simply not possible). Much of the experience of this community is embodied in several generations of systems, from SIMNET to NPSNET-IV [13] and PaRADISE [15], and in two generations of standards: Distributed Interactive Simulation (DIS)[1] [10] and High-Level Architecture (HLA) [11]. The community has its own conferences (e.g. Simulation Interoperability Workshops), publication channels and user groups, which are generally very active and often government managed.

The goals of this community are to enable the staging of large, high-fidelity, relatively inexpensive distributed training and simulation exercises.

The ideology of this group has the following key components:

- It has to work in the near future (not necessarily today), using available mid-range technologies.
- Dedicated bandwidth and networks are available (and may be required for security reasons, if nothing else).
- It has to work for a large number of simultaneous users within a single exercise (100s today, 1000s or 100,000s soon).
- User activities within the worlds are highly constrained, e.g. driving a tank, firing shells, being blown up.
- The user group is well-defined, available for training, and can be assumed to have a certain level of expertese.
- It must be interoperable, at least in deployment situations.
- It must be effective for training.

These factors have encouraged approaches optimized to support constrained (often vehicle-based) interaction between several hundred entities (users or simulations) on a dedicated (or at least high-bandwidth) network, using broadcast or multicast network protocols. Partly because of the use of broadcast/multicast, and partly because of the application domain, the common conception of distribution in this domain is of dynamic publishing of information about autonomous user movements and activities (such as firing).

Strengths of this community's work include extensive consideration of scalability with regard to the number of simultaneous users (e.g. through interest management), and leading the use of multicast communication. On the other hand, there has been only limited support given for broader ranges of use, and for less capable (or highly variable) infrastructures.

4 GAMES

Large segments of the game writing community are looking to multi-played gaming as a major strategic area for current and near-term development. Examples include a plethora of networked first-person shooters (i.e. Doom, Quake, et al.), and multi-user role-playing games (including Ultima Online [1]). The gaming community is oriented around shows more than conferences or journals (e.g. ECTS).

The prime objective of this group is to sell product, whether software or playing time (or support). To date, this is the only segment of the multi-user virtual world community that has to show a direct profit right now.

The ideology of this group has the following key components:

- It must work right now, on the machines and networks that people already have (like Internet virtual worlds).
- Each product exists largely in isolation: it can be innovative, and completely incompatible with any other system or release.
- There is intense competition and rapid innovation.
- Traditionally, there has been no ongoing support, maintenance or development, although this model breaks down with long-running online experiences such as Ultima Online.
- The end-user experience is paramount.
- Typically, networking is not a core company competence, and is an "add-on" for an initially stand-alone product type.
- The end-user cannot be trusted, and may try to cheat in quite sophisticated ways.

These factors lead to a fragmented community, divided by concerns for secrecy, market lead and share. There is no culture (and no motive) for sharing technical knowledge or innovation, although there may be "learning by inspection" and/or reverse engineering of other products. Solutions are custom-built, often around client-server architectures (which are simple and reliably deployable).

[1] The Web3D dis-java-vrml working group is providing a VRML/Java-based implementation of DIS. The vrtp working group's brief is larger, but its model of networking and distribution is currently closest to this community's.

Strengths of this community's work are short time-lapse between successive generations of system and great scope to innovate. The hard fiscal constraints also encourage a very realistic approach. On the other hand, this is a very difficult community to integrate with the broader community, or to share knowledge with (other than one-way). Relatively recently a number of companies have offered "middleware" solutions for networking targeted specifically at the games community (e.g. RTime [16]). These products are subject to ongoing development, and greater investment in time and expertese than any single game company would typically have available for networking. However they are still commercial products and have little ongoing interaction with the broader research community.

5 VISUALIZATION

Visualization is a primary application for 3D computer graphics, across a range of disciplines. Types of visualization include architectural (of buildings, built or otherwise), scientific (of scientific data and simulations), environmental (of the natural world and sensor data), engineering (of planned products and designs), and information (of abstract data). The majority of visualization work is for single user systems, however a number of groups are extending these visualizations to support multiple users, for example EVL [17]. This community tends to publish and gather around established visualization and high-performance computing journals and conferences (e.g. International Symposium on High Performance and Distributed Computing, IEEE Visualizaton).

The primary objective of this community is to extend existing approaches to visualization to make them accessible to multiple concurrent (and/or asynchronous) users. This is either in support of a specific existing user community, or within the visualization research community.

The ideology of this group has the following key components:

- The visualization itself is critical.
- Even for a single user, the visualization typically requires high-end machines to generate and interact with.
- Consequently user machines are generally state of the art, and networks are typically very high bandwidth.
- Funding is typically from government or similar sources, and outputs are knowledge and/or abstract value rather than fiscal.
- The number of simultaneous users is very small (2 would be a good place to start).
- Projects are either for "real" – often in-house – users, or are proof-of-concept or motivating applications for advanced network and computing infrastructures.

As noted, the visualization itself is typically much larger (in data) and more demanding than the multi-user aspects of the problem within this domain. Much effort has to be directed towards distributing the visualization itself.

Strengths of the community are its addressing of very large, data intensive virtual worlds (the visualizations themselves), and their access to – and proving of – next-generation computers and networks. On the other hand, many of the issues which concern other groups (like supporting large numbers of users, or access over dial-up networks) are non-issues, and consequently ignored.

6 "PURE" RESEARCH

By "pure" research, I mean that this community is typically concerned with multi-user virtual reality technology for its own sake. It doesn't typically matter if there are no users: it is the concept that matters. Subsidiary interests range from the users' experiences to the network protocols used. Examples of work in this area include our own (e.g. the MASSIVE systems [8]), SICS DIVE [9], Broll [7] (unusual in having a strong VRML emphasis) and a rich history of prior systems such as VEOS [6], BrickNet [14], etc.

The primary objective of this community is to publish conference and journal papers about new techniques and/or evaluations of existing techniques. Typical publication venues include well-known conferences (IEEE VR, ACM VRST, CHI, CSCW) and journals (Presence, Multimedia, Computer Graphics and Applications). There are also a small number of companies that are or have been active within the research community, but have also delivered their research as products or pre-products under what I would class the Internet virtual worlds category; these include MERL (SPLINE [5] and Open Community [4]) and Sony (Community Place [12]).

The ideology of this group has the following key components:

- Publications, and to a lesser extent demonstrations, are what really matters (excepting the commercial groups, noted above).
- Available networks and technologies cover a broad spectrum, from dedicated research networks and machines to commodity PCs and networks (although there is typically less support for domestic-type dial-up users).
- There is limited continuity of development.
- Systems may be very focused, and limited in breadth (addressing a single research topic).
- The majority of work is home-grown, with little sharing of resources beyond certain common facilities (e.g. graphics, device handling).
- Access to source code is always required, making the use of most commercial products difficult or impossible.
- The competitive publishing culture requires that work be related to prior knowledge, but does not necessarily encourage groups to take on and advance the best parts of other's work (it may be hard to publish a paper that just agrees with an existing publication).
- There is only limited motivation to exploit research, or to ensure its effective transfer beyond this single community.

These factors give rise to a proliferation of independent, often partial systems, which evolve and die over relatively short time scales. There is relatively little exploitation of common systems and technologies, and a significant amount of re-engineering (to give the greatest possible flexibility and/or control).

Strengths of this community are its addressing of a very broad range of technical issues relating to multi-user virtual worlds, the constant innovation and introduction of new techniques and approaches, and the availability of information, through publications. On the down side, work can be cut off from today's pragmatic concerns – systems may only ever be used by the people who build them – and software is often unsuited for outside exploitation (at least without significant further

development). It can also be very hard to integrate or combine work from different groups, which often rests on significantly different philosophies and assumptions.

7 A TAXONOMY OF APPROACHES TO CREATING VIRTUAL WORLDS

While considering these various interests and approaches I was forced to consider whether there were any *fundamental* differences between the technical approaches adopted. This led me to articulate the following taxonomy of approaches to creating multiuser virtual worlds, in an attempt to clarify the differences.

It is apparent that different systems and approaches have significantly different capabilities and ambitions. I have tried to identify the critical levels of distribution support that can be found in different systems; this gives a more technically-oriented perspective on some of the differences between the communities described above. Five levels of multi-user support are listed in order below, from the simplest to the most comprehensive.

1. Shared names for virtual worlds, for example URLs for single-user VRML worlds. This is not actually multi-user, but is a necessary pre-requisite. For example, how do we know that we are or were in the same world? Because it has the same name – without a common method for identifying worlds we will be unable to meet.
2. Shared representations of presence and identity, i.e. mutually visible embodiments or avatars. These represent users to themselves and to one another within a shared world. The simplest shared virtual world to build is just a number of copies of the world plus a hand-crafted mechanism to add avatars. This may be supplemented with some kind of additional communication facility, such as text or audio chat. Simple visualization applications can exist quite happily at this level. This is typical of students' first attempts at a multi-user VR system.
3. Shared behavior triggers, i.e. certain key actions by one user will be visible to other users (e.g. opening the door, triggering a particular animation, firing munitions). Philosophically, this is the state of most VRML multi-user systems: distribution is added on a case-by-case, object-by-object basis to achieve an adequate agreement about what is "happening". It is also typical of most distributed simulation systems in practice (although some are capable of more).
4. Sharing arbitrary state and updates, i.e. the world (or portions of it) can be changed in any way, by any user, and everyone will see "the same thing" (subject to certain constraints, e.g. time delays). "Pure" research systems typically exist around this level.
5. Sharing arbitrary physics and constraints, i.e. not just what the world looks like and is doing, but what is *should* do under any set of circumstances. This is an extremely hard problem in the general case, because of severe problems coordinating processing and decisions across networked processes engaged in close interaction.

These levels are shown diagrammatically in figure 1.

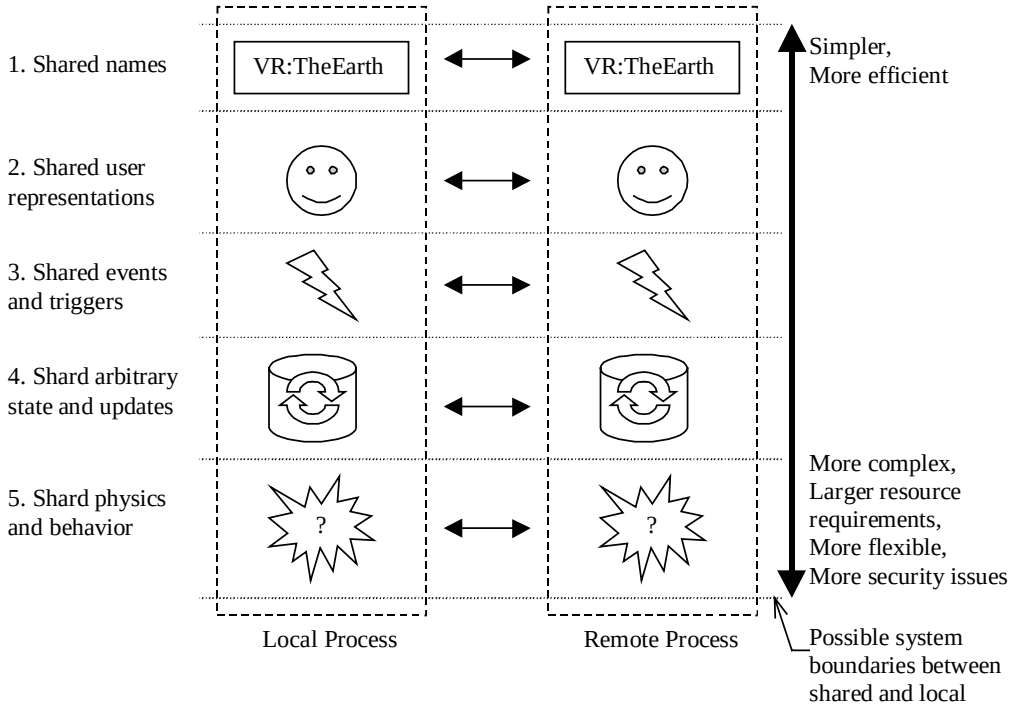

Figure 1. Levels of multi-user support

7.1 General Observations

In general, the earlier approaches are simpler to do, more efficient, require fewer resources (computing and networking), and have more limited security implications. However they are also less flexible and capable than the later approaches: there are some things that they simply cannot do (and other things that they cannot do simply). Of the communities described those with significant pragmatic constraints can be viewed as starting at point 1 and adding networked capabilities only as necessary. Addressing their core problems (low bandwidth, large visualization, rapid shrink-wrapped product development, etc.)

means that they do not have the resources (or the motivation) to be more ambitious than necessary.

Those with fewer constraints have the liberty to shoot for higher target from the outset; indeed, they may not even need to be successful (i.e. have a working system) in order to succeed (i.e. publish a paper). Even when they do produce a working system or useful technique it often requires resources or assumptions which are unsupportable in the other communities (e.g. congestion-free networks).

7.2 Using the Taxonomy

I am not suggesting that everyone must or should aim for level 5, although there is a general upward trend. For example, consumer technology has reached (or is rapidly approaching) the point at which effective networked 3D virtual worlds can be supported. This, plus ever-growing user expectations, is forcing many groups to head further up the scale of multi-user support. More generally, at any given level of support some things will be easy and natural to do, others will be possible but hard, and some things will simply be impossible. When building multi-user virtual worlds it is necessary to consider what level of capability is actually required, because each imposes its own costs and requirements.

The taxonomy also illustrates one way in which our various communities are converging, pushed in part by the continuous development of computer and network technology, and the maturing of multi-user virtual worlds as a research area.

On the one hand, the gap between "research" and "consumer" computing and networking technologies appears to be narrowing rapidly. Communities oriented towards consumer applications are able to push up the scale as entry-level computer and networking technology improves.

On the other hand, from the more research oriented communities there is a growing interest in practical deployment issues, such as supporting varying infrastructures, providing managed qualities of service, and integrating security and trust considerations. This work will allow these technologies to be deployed in contexts for which they have so far been unsuited. These may be viewed as complementary bottom-up and top-down approaches leading towards a more generally applicable multi-user virtual world framework or solution set.

8 CONCLUSIONS

As the area of multi-user virtual worlds matures the "entry cost" continues to rise, i.e. the definition of a minimum acceptable system escalates continuously, driven by people's growing experience and appreciation of the possibilities. This either requires continuously increasing investment to catch up, or greater communication, collaboration, and sharing of resources, techniques and software. For several years the research community has been discussing a more effective day-to-day collaboration; it appears that this will soon be unavoidable for many. Today, to claim to be comprehensive, a multi-user virtual world system must address all of the following areas:

- Consistency.
- Behavior, virtual "physics".
- Scalability, of users, worlds and data.
- Heterogeneity of deployment platforms and networks.
- Performance and efficiency.
- Interest management (for scalability, heterogeneity, efficiency, etc.).
- Quality of Service issues.
- Persistence.
- Security, trust, access and governance.
- Interaction, and interactivity.
- Navigation.
- Support for collaboration.
- Content creation and management.
- Versioning.

This list, and our appreciation of the complexities of each area, is only likely to grow.

8.1 Next Steps

I am very reluctant to suggest where we should go next; I do not feel that I have reached the necessary comprehensive understanding of multi-user VR, its possibilities and its technical requirements and tradeoffs (I don't believe anyone has, yet). However I make the following observations regarding work within the context of VRML.

- The definition of a reduced core for VRML in the X3D work, plus the scope to define multiple co-existing profiles (including for multi-user support) seems to me more appropriate than the ambition to create a single solution, even now.
- The same data types and update mechanisms which will support streaming of updates into virtual worlds should underlie at least one model of generic multi-user support.
- One of these models should be a distributed scene-graph (multi-user support at level 4 in the taxonomy), with the facility to introduce non-distributed sub-graphs in a principled manner.
- Single-user VRML should be considered a special case of multi-user, and expressed accordingly. This could support more varied models of distribution.
- The representation of the local user should be detached from the browser representation, so that the same representation can be used for remote users.
- Behaviors should then be expressed in terms of these user representations, for example proximity sensors should respond to the proximity of one or more users (local and/or remote), and be able to represent this.

Of course, these suggestions reflect my own interests and prejudices…

8.2 Final Words

As our aspirations grow, and as our individual "hard problems" become a little more tractable, we each have the motivation and opportunity to grow our common understanding and the common ground between our respective communities. One hope for X3D is that it will release VRML from its traditional "one size fits all" straightjacket. This is particularly important in the area of multi-user support, because there are so many different ideologies in play: it is naïve to expect a universal solution within the next five years (if ever).

Acknowledgements

I would like to thank all those who have participated in the System Aspects of Sharing a Virtual Environment workshop series (http://vr.edu/capps/sharedvr), for the insights and concerns they have shared from their own communities. I also acknowledge the support of the UK's EPSRC through project grants GR/M 21256, GR/L02074 and GR/M09223.

References

[1] Electronic Arts, Ultima Online, http://www.owo.com/ (verified Sept. 23rd 1999).

[2] Blaxxun Interactive, http://www.blaxxun.com/ (verified Sept. 24th 1999).

[3] VR Telecom, http://www.vrtelecom.com/ (verified Sept. 24th 1999).

[4] Mitsubishi Electric ITA, Open Community, http://www.opencommunity.com/ (verified Sept. 24th 1999).

[5] Barrus, J.W., Waters, R.C, and Anderson, D.B., "Locales: Supporting Large Multiuser Virtual Environments", in *IEEE Computer Graphics and Applications*, Nov. 1996, 16(6), pp. 50-57.

[6] Bricken, William, and Coco, Geoffrey, "The VEOS Project", *Presence*, 3(2), Spring 1994, pp. 111-129.

[7] Broll, W., "Distributed Virtual Reality for Everyone – a Framework for Networked VR on the Internet", *IEEE VRAIS'97*, March 1-5, Albuquerque, NM.

[8] Greenhalgh, C. M. and Benford, S. D., Supporting Rich And Dynamic Communication In Large Scale Collaborative Virtual Environments, *Presence: Teleoperators and Virtual Environments*, Vol. 8, No. 1, February 1999, pp. 14-35, MIT Press.

[9] Hagsand, Olof, "Interactive Multiuser VEs in the DIVE System", *IEEE Multimedia*, Spring 1996, 3(1), pp 30-39, IEEE Computer Society, ISSN 1070-986X

[10] IEEE, International Standard, ANSA/IEEE Std 1278-1993, Standard for Information Technology, Protocols for Distributed Interactive Simulation (March 1993).

[11] DMSO High Level Architecture home page, http://hla.dmso.mil/ (verified Sept. 24th 1999).

[12] Lea, Rodger, Honda, Yasuaki, and Matsuda, Kouichi, "Virtual Society: Collaboration in 3D Spaces on the Internet", *Computer Supported Cooperative Work: The Journal of Collaborative Computing*, 6:227-250, 1997, Kluwer Academic Publishers: Netherlands.

[13] Macedonia, Michael R., Zyda, Michael J., Pratt, David R., Brutzman, Donald P., and Barham, Paul T., "Exploiting Reality with Multicast Groups: A Network Architecture for Large-Scale Virtual Environments", *Proc. 19995 IEEE Virtual Reality Annual International Symposium (VRAIS'95)*, 11-15 March, 1995, RTP, North Carolina.

[14] Singh, Gurminder, Serra, Luis, Pang, Willie, and Ng, Hern, "BrickNet: A Software Toolkit for Network-Based Virtual Worlds", *Presence*, 3(1), Winter 1994, pp. 19-34.

[15] Singhal, S.K., and Cheriton, D.R, "Using Projection Aggregations to Support Scalability in Distributed Simulation", in *Proc. 1996 International Conference on Distributed Computing Systems (ICDCS'96)*.

[16] RTIME Inc., http://www.rtimeinc.com/ (verified Sept. 24th 1999).

[17] Leigh, J., Johnson, A., DeFanti, T., Bailey, S., Grossman, R., "A Methodology for Supporting Collaborative Exploratory Analysis of Massive Data Sets in Tele-Immersive Environments", in the *Proceedings of the 8th IEEE International Symposium on High Performance and Distributed Computing*, Redundo Beach, California, Aug 3-6, 1999.

Biographical Note

Chris Greenhalgh is a lecturer in Computer Science at the University of Nottingham, UK, where he gained his PhD in 1997, based on his work on large scale collaborative virtual environments. His thesis was a winner of the 1998 UK Distinguished Dissertations in Computer Science competition, and has been published by Springer. He has been designing and building multi-user virtual reality systems for six years and is the primary architect of the MASSIVE-1, MASSIVE-2 and MASSIVE-3 systems.