Proc. of First IEEE Symposium on Information Visualization, Atlanta, GA, Oct. 1995, pages 97–104.

Case Study

Visualising Cyberspace: Information Visualisation in the Harmony Internet Browser

Keith Andrews

Institute for Information Processing and Computer Supported New Media (IICM)
Graz University of Technology,
A-8010 Graz, Austria.

kandrews@iicm.tu-graz.ac.at

Abstract

The explosive growth of information systems on the Internet has clearly demonstrated the need to organise, filter, and present information in ways which allow users to cope with the sheer quantities of information available.

The scope for visualisation of Gopher and WWW spaces is restricted by the limitations of their respective data models. The far richer data model supported by the Hyper-G Internet information system is exploited by its Harmony client to provide a number of tightly-coupled, two- and three-dimensional visualisation and navigational facilities, which help provide location feedback and alleviate user disorientation.

1 Introduction

There has been strong recent interest in the application of visualisation techniques to network information systems. The explosive growth of information systems on the Internet has clearly demonstrated the need to organise, filter, and present information in a way which allows users to cope with the sheer quantities of information available. More traditional (and arguably more functional) two-dimensional visualisations have been pushed into the sidelines as glamorous three-dimensional visualisations take the limelight. This paper presents the tightly-coupled, two and three-dimensional visualisations implemented in the Harmony client for the Hyper-G Internet information management system.

First let me distinguish between *hand-crafted* and *automatically generated* presentation of information. Hand-crafted presentation entails careful, prior construction of an information layout by a designer using appropriate tools. Such visualisations are generally designed for very specific applications and are often quite intricate in nature. Automatically generated visualisations are composed on-the-fly by the underlying system according to a set of pre-supplied layout rules and components. A neutral file format like the emerging Virtual Reality Modeling Language (VRML) [1] for 3D models, allows the distinction to be blurred somewhat at least as far as final presentation is concerned, since both hand-crafted and automatically generated VRML files can be displayed by one and the same viewer.

One issue affecting automatically generated visualisations is that of dynamic restructuring, for example to take into account changes in the database or in response to user queries or filters. In some cases this may be desirable, but in many cases an inherent advantage of spatial visualisations, that users can incrementally build a mental map of places and locations and navigate by visual memory, will be lost when the information space continually changes its form. At the very least, such changes should be made via smooth transitions as far as possible so as to take advantage human visual perception and object constancy [2].

Other issues include the choice of visual representations and navigational metaphors. What kind of visual representations are best? Two or three-dimensional icons of books for text documents, framed pictures for images, radios emitting audio clips, TV sets for video, etc. Should planar or truly 3D structure maps be used, indoor or outdoor scenes? In 3D visualisations, typical navigational metaphors include walking, flying, driving, teleportation, and magic carpets. However, why not bring the information to the user, rather than having the user navigate to the information? Indeed, the use of user-configurable home spaces, where users can place and organise frequently accessed information (or handles to it) might reduce the need for navigation in the first place. Whilst ideas abound and numerous techniques have been implemented, there has upto now been surprisingly little usability testing of information visualisation techniques.

2 Related Work

Two-dimensional visualisations of information structures are well covered in the literature. A good overview of 2D link maps in hypermedia systems can be found in [3], fisheye viewing techniques for graphs are detailed in [4]. Space-filling visualisations of hierarchically structured information using tree maps are described in [5], feature maps of arbitrary document spaces are described in [6], and hyperbolic visualisations of large trees in [7], to mention just a few techniques.

Perhaps the most well-known work in the area of three-dimensional information visualisation is that of Card et al [2, 8] at Xerox PARC on the Information Visualiser and 3D/Rooms. They allow users to interactively explore workspaces modeled as three-dimensional rooms. Particular data sets (nodes) are visualised in rooms of their own, and doors lead from one room to another in a manner similar to hypermedia links. The Information Visualizer provides three-dimensional representations for linear and hierarchically structured information: the perspective wall [9] and cone tree [10] respectively. Linear information, such as chronologically ordered information, is pasted on to a virtual wall from left to right. The wall has a large front section, and left and right sides which tail off into the background. Information can be slid along the wall to bring it into focus on the front section. The information can also be stretched or shrunk along the wall. Hierarchical information, such as part of a file system or a company hierarchy chart, is laid out as a uniform 3D cone. The tree can be rotated to bring interesting parts to the front and pruned to remove non-relevant information.

VizNet [11] also uses a cone tree representation for hierarchical information, but provides an additional spherical representation for associative relationships (local map). The current node is located (say) at the north pole, nodes similar to it are strung along lines of longitude. Lower level objects are displayed on lower level spheres (like peeling away layers of an onion).

SemNet [12] was an exploratory system which represented knowledge bases as directed graphs in 3D. Labeled rectangles (nodes) were connected by lines or arcs. The 3D layout has the advantage over 2D layouts that the nodes of an arbitrary graph can be positioned so that no arcs intersect. Several techniques for positioning nodes were explored: random, multidimensional scaling, heuristics, and manual editing. Clustering techniques and fisheye views were also implemented.

Serra et al [13] discuss the use of 3D object hierarchies with attached multimedia documents. Each component in the 3D object hierarchy (part-of relationships) may be combined into a *concept node* with text, image, and video documents. Links may be made from these text, image, and video documents to other concept nodes. However, there is no support for arbitrary links from (parts of) the 3D object as such, and the 3D object hierarchy itself forms the entire extent of the hyperstructure.

Smith and Wilson [14] describe a prototype system based on HyperCard and Virtus Walkthrough (a 3D visualisation system), in the context of an academic departmental information system. They enabled users to interactively explore a 3D model of the department: when they approached within a certain distance of a source anchor, it automatically triggered to display a corresponding text document.

The File System Navigator (FSN, or "Fusion") written by Joel Tesler and Steve Strasnick at Silicon Graphics [15] visualises a Unix file system as an information landscape. Directories are represented by blocks laid out on a plane, their height representing the cumulative size of the contained files. Smaller blocks atop the directory blocks represent files in the directory (their size also mapped to their height). Users can "fly" over the landscape, taking it in as a whole, or swoop down to a specific directory. Clicking on the arc to a subdirectory results in an invisible hand grabbing you and leading you through space to that subdirectory. Clicking on a file block brings a virtual spotlight to bear on that block, double-clicking opens the file for editing, etc.

The above systems are for the most part proof-of-concept prototypes for individual visualisation techniques. In Graz, we have integrated several tightly-coupled information visualisation techniques (including a two-dimensional graphical structure map generator, a three-dimensional document browser and a three-dimensional information landscape) into a real, day-to-day information tool running over the Internet: the Harmony client for the Hyper-G hypermedia information system [16, 17].

Visualisation techniques are being applied to other Internet information systems such as Gopher [18] and the World-Wide Web (WWW) [19]. GopherVR [20] is an experimental 3D spatial interface to the hierarchically structured information on a Gopher server. Visualisations of WWW spaces such as Webspace [21], the Navigational View Builder [22], and HyperSpace [23] have begun to appear. However, the scope for visualisation of Gopher and WWW spaces is severely restricted by the limitations of their respective data models: Gopher has only hierarchical structures and WWW has only forward hyperlinks. Hyper-G, on the other hand, has a much richer data model on which to base visualisations: a combination of hierarchical structure, (bidirectional) hyperlinks, and fully integrated search and retrieval facilities. The rest of this paper looks at the information visualisation facilities as currently implemented in the Harmony client for Hyper-G.

3 The Hyper-G Data Model

Information in Hyper-G may be structured both hierarchically into so-called *collections*, and by means of associative hyperlinks. A special kind of collection called a *cluster* groups logically related or multilingual versions of documents. Every document and collection must belong to at least one collection, but may belong to several. Navigation may be performed down through the collection hierarchy (the collection "hierarchy" is, strictly speaking, a directed acyclic graph), access rights assigned on a collection-by-collection basis, and the scope of searches restricted to particular sets of collections. Collections may span multiple Hyper-G servers, providing a unified view of distributed resources.

Links in Hyper-G are stored in a separate link database and are bidirectional (directed, but may be followed backwards): both the incoming *and* outgoing hyperlinks of a document are always known and available for visualisation. Furthermore, Hyper-G has fully integrated search facilities including full text search with relevance scores and some limited support for similarity measures between documents.

All in all, the richness of the Hyper-G data model provides plenty of scope upon which to base visualisations: hierarchical structure, (bidirectional) hyperlinks, and search and retrieval facilities. The Harmony client for Hyper-G exploits this richness to provide tightly-coupled two- and three-dimensional visualisation and navigational facilities help provide location feedback and alleviate user disorientation.

4 Dynamic 2D Structure Maps

The Harmony Session Manager (see Figure 1, top left) provides navigation through the collection structure, search facilities, and various general functions such as user identification and language selection. Collections may be opened and closed and clusters or individual documents activated within the graphical collection display by double-clicking. Central to the design of Harmony is the concept of *location feedback*. When a document or collection is visited, its location within the collection structure is automatically displayed in the Session Manager's collection browser (by opening up the path to it), regardless of whether the object was reached as the result of a search, by following a hyperlink, or via the local map. This simple technique is a powerful instrument in the fight against becoming "lost in hyperspace" – users can orient themselves with reference to a fixed structural framework. In the case of search results and the local map, mere selection of an object initiates location feedback, providing users with a sense of the context of an object, prior to any decision to view it (see Figure 1).

The Harmony Local Map facility provides a kind of short-range radar, generating on request (dynamically) a map of the hyperlink neighbourhood of a chosen document, similar to the local map of Intermedia [3]. However, Harmony's local map can also show other relationships, such as collection membership, annotations, inline images, and the textures applied to a 3D model, as can be seen in Figure 2. The number and type of incoming and outgoing relationships can be interactively configured in the Options panel. Single-clicking selects an object and activates location feedback, displaying the object's location in the collection structure. Users can navigate within the local map by generating a new map around any selected object; objects are accessed by double-clicking. One practical use of the Local Map is to check whether a particular inline image is used in multiple text documents before changing or deleting the image.

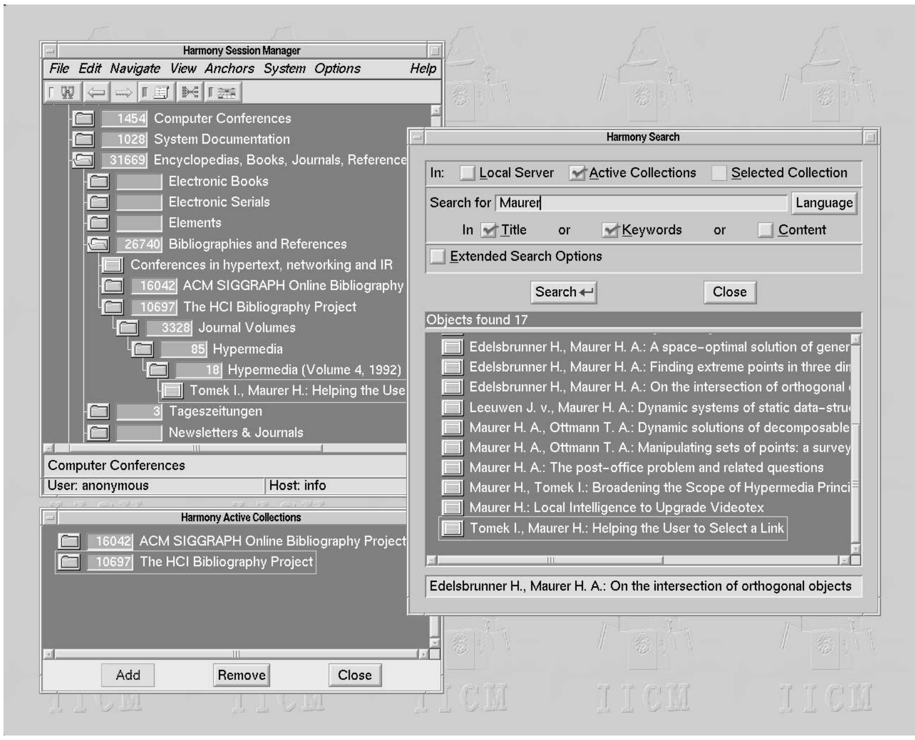

Figure 1: Location Feedback in Harmony

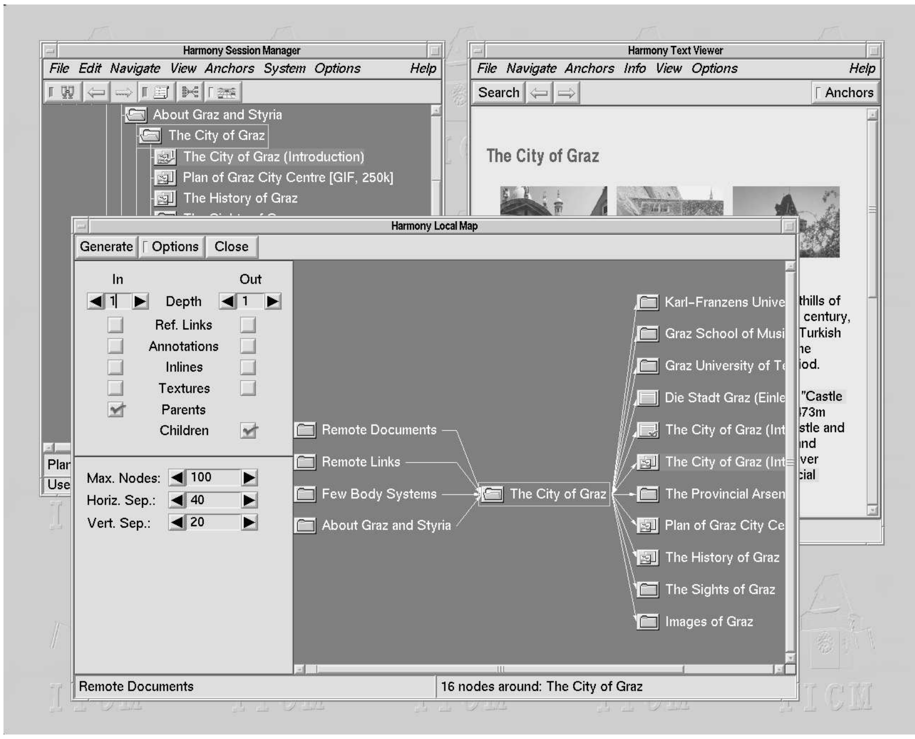

Figure 2: The Harmony Local Map

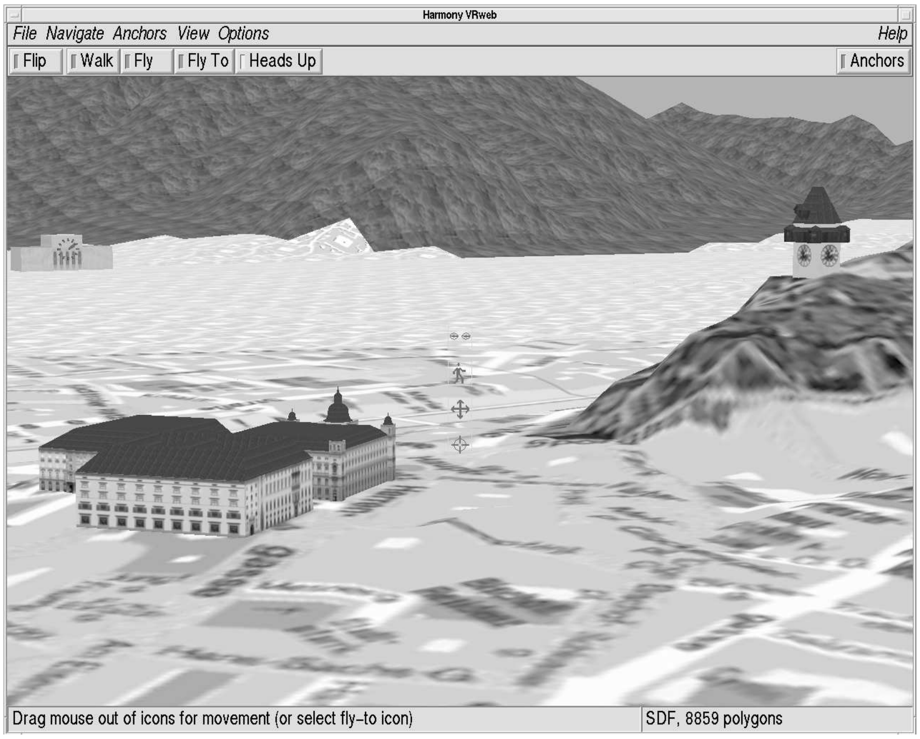

Figure 3: Harmony's VRweb 3D Viewer

5 3D Hyperdocuments

The Harmony VRweb 3D scene viewer displays model description files representing arbitrarily complex objects or scenes. The models may either be of the hand-crafted or of the automatically generated variety, depending upon the application. Both VRML and SDF, a file format derived from the Wavefront Advanced Visualizer, are supported. In Figure 3, a hand-crafted 3D plan of the centre of Graz features some of the sights of the city; embedded hyperlinks lead to further information about the various sights. Users typically explore a model of a scene by moving themselves (walk, fly, fly to, and heads-up navigation modes) and view a model of an object by moving the model (translate, rotate, zoom).

6 3D Structure Maps

Harmony's Information Landscape, shown in Figure 4, is an interactive, three-dimensional visualisation of the collection structure, tightly coupled to the Session Manager's two-dimensional collection browser display (changes in one are reflected in the other). The collection hierarchy is mapped out onto a plane, documents within a collection are arranged on top of the corresponding block; colour and height are currently used to encode document type and size respectively. Users can "fly" over the landscape looking for salient features, like flying over a file system directories with FSN [15]. A flat overview window (top right) provides a further aid to orientation. Whereas FSN reads in and lays out the entire directory hierarchy once at program start-up, Harmony's Information Landscape is constructed and laid out incrementally as users open and close collections and subcollections, access documents, and perform searches.

Work on the Information Landscape is still at an early stage. We are currently experimenting with the use of textured landscapes (as shown in Figure 5), and will introduce simple 3D icons to replace the basic blocks currently used. A default set of 3D icons will represent document type, with colouring used to represent document age and hence giving a rapid visual impression of where new information is located. The author of a collection will be able to specify an arbitrary model as its 3D icon: for example a model of the Eiffel Tower to represent a collection about Paris. A further innovation will be the superimposition of

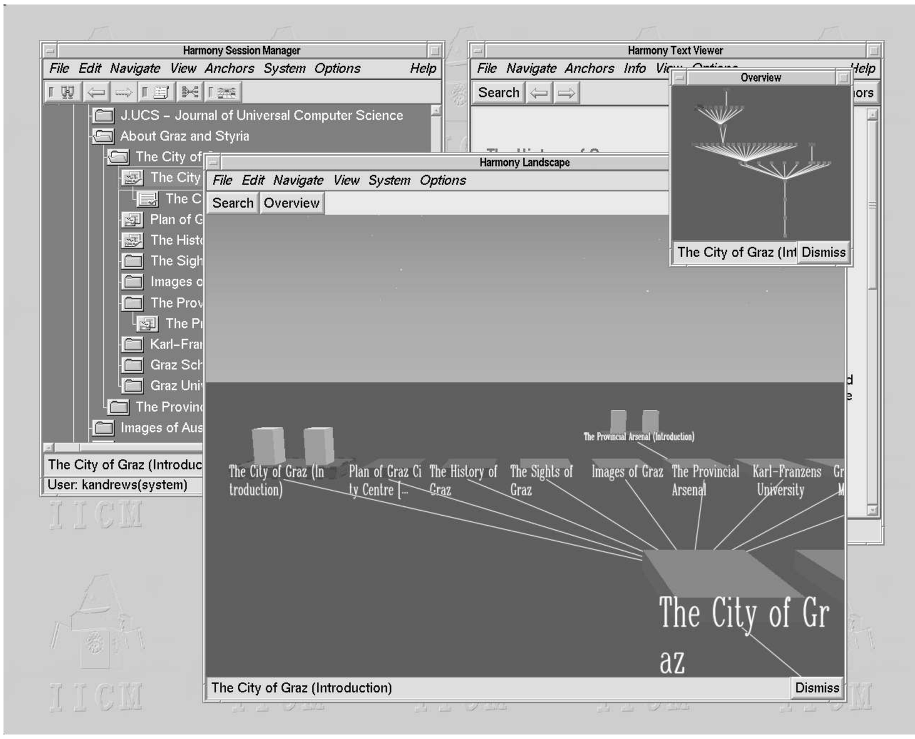

Figure 4: The Harmony Information Landscape

Figure 5: Textured Information Landscape

arced cables upon the landscape to represent the link relationships of a specific document. Through their ability to compactly display many thousands of objects, 3D visualisations are perhaps the only effective means of browsing in and judging the extent of large, dynamic information spaces.

7 Concluding Remarks

Following an overview of information visualisation techniques and current efforts to visualise Internet information spaces, I have presented in more detail the visualisation facilities currently available in the Harmony client for Hyper-G. Recent informal usability evaluations of Harmony (thinking aloud tests to provide design feedback) have shown in particular that the close coupling of several orthogonal visualisation and navigational facilities in Harmony helps provide location feedback and alleviate user disorientation.

In the near future, both hand-crafted and dynamically generated visualisations will start to abound, as modeling languages like VRML take hold and appropriate software becomes available for a wide variety of applications. The announcement of VRML browsers for Gopher, WWW, and Hyper-G can be expected to rapidly increase the number of visualisations of these information spaces.

Further information about Hyper-G and Harmony and installation details may be retrieved by anonymous ftp from `ftp.iicm.tu-graz.ac.at` in directory `/pub/Hyper-G` or from the IICM Information Server under `http://info.iicm.tu-graz.ac.at/` or `gopher://info.iicm.tu-graz.ac.at/`.

8 Acknowledgements

The financial support of Hyper-G by the Austrian Ministry of Science, JOANNEUM RESEARCH, and the European Space Agency is gratefully acknowledged. Current work on Virtual Information Spaces is partly financed by the Anniversary Fund of the Austrian National Bank, under project number 5334.

I am indebted to the members of my Harmony project team for their valuable input, hard work, and dedication.

References

[1] Gavin Bell, Anthony Parisi, and Mark Pesce. The Virtual Reality Modeling Language – Version 1.0 Specification. Available at `http://vrml.wired.com/vrml.tech/vrml10-3.html`.

[2] George G. Robertson, Stuart K. Card, and Jock D. Mackinlay. Information Visualization Using 3D Interactive Animation. *Communications of the ACM*, 36(4):56–71, April 1993.

[3] Kenneth Utting and Nicole Yankelovich. Context and Orientation in Hypermedia Networks. *ACM Transactions on Information Systems*, 7(1):58–84, January 1989.

[4] Manojit Sarkar and Marc H. Brown. Graphical Fisheye Views of Graphs. In *Proc. CHI'92*, pages 83–91, Monterey, May 1992. ACM.

[5] Brian Johnson and Ben Shneiderman. Tree-Maps: A Space-Filling Approach to the Visualization of Hierarchical Information Structures. In *Proc. IEEE Visualization '91, San Diego, CA*, pages 284–291, Los Alamitos, CA, October 1991. IEEE Computer Society Press.

[6] Xia Lin. Visualization for the Document Space. In *Proc. IEEE Visualization '92, Boston, MA*, pages 274–281, Los Alamitos, CA, October 1992. IEEE Computer Society Press.

[7] John Lamping and Ramana Rao. Laying out and Visualizing Large Trees Using a Hyperbolic Space. In *Proc. UIST'94, Marina del Rey, CA*, pages 13–14, New York, November 1994. ACM.

[8] Stuart K. Card, George G. Robertson, and Jock D. Mackinlay. The Information Visualiser, An Information Workspace. In *Proc. CHI '91: Human Factors in Computing Systems*, pages 181–188, New York, May 1991. ACM.

[9] Jock D. Mackinlay, George G. Robertson, and Stuart K. Card. The Perspective Wall: Detail and Context Smoothly Integrated. In *Proc. CHI'91, New Orleans, Louisiana*, pages 173–179, New York, May 1991. ACM.

[10] George G. Robertson, Jock D. Mackinlay, and Stuart K. Card. Cone Trees: Animated 3D Visualizations of Hierarchical Information. In *Proc. CHI'91, New Orleans, Louisiana*, pages 189–194, New York, May 1991. ACM.

[11] Kim Michael Fairchild, Luis Serra, Ng Hern, Lee Beng Hai, and Ang Tin Leong. Dynamic FishEye Information Visualizations. In Rae A. Earnshaw, Michael A. Gigante, and Huw Jones, editors, *Virtual Reality Systems*, pages 161–177. Academic Press, 1993.

[12] Kim Michael Fairchild, Steven E. Poltrock, and George W. Furnas. SemNet: Three-Dimensional Representations of Large Knowledge Bases. In Raymonde

Guindon, editor, *Cognitive Science and its Applications for Human-Computer Interaction*, pages 201–233. Lawrence Erlbaum, Hillsdale, NJ, 1988.

[13] Luis Serra, Tat-Seng Chua, and Wei-Shoong Teh. A Model for Integrating Multimedia Information Around 3D Graphics Hierarchies. *The Visual Computer*, 7(5-6):326–343, May/June 1991.

[14] Pauline A. Smith and John R. Wilson. Navigating in Hypertext Through Virtual Environments. *Applied Ergonomics*, 24(4):271–278, August 1993. Butterworth-Heinemann Ltd.

[15] Joel Tesler and Steve Strasnick. *FSN: The 3D File System Navigator*. Silicon Graphics, Inc., Mountain View, CA, 1992. Available by anonymous ftp from `sgi.sgi.com` in directory `sgi/fsn`.

[16] Keith Andrews, Frank Kappe, and Hermann Maurer. Hyper-G: Towards the Next Generation of Network Information Technology. *Journal of Universal Computer Science*, 1(4):206–220, April 1995. Special Issue: Selected Proceedings of the Workshop on Distributed Multimedia Systems, Graz, Austria, Nov. 1994. Available at `http://info.iicm.tu-graz.ac.at/Cjucs_root`.

[17] Keith Andrews, Frank Kappe, and Hermann Maurer. Serving Information to the Web with Hyper-G. *Computer Networks and ISDN Systems*, 27(6):919–926, April 1995. Proc. Third International World-Wide Web Conference, WWW'95, Darmstadt, Germany.

[18] Mark P. McCahill and Farhad X. Anklesaria. Evolution of Internet Gopher. *Journal of Universal Computer Science*, 1(4):235–246, April 1995. Selected Proceedings of the Workshop on Distributed Multimedia Systems, Graz, Austria, Nov. 1994. Available at `http://info.iicm.tu-graz.ac.at/Cjucs_root`.

[19] Tim Berners-Lee, Robert Cailliau, Ari Luotonen, Henrik Frystyk Nielsen, and Arthur Secret. The World-Wide Web. *Communications of the ACM*, 37(8):76–82, August 1994.

[20] Mark P. McCahill and Thomas Erickson. Design for a 3D Spatial User Interface for Internet Gopher. In *Proc. of ED-MEDIA 95*, pages 39–44, Graz, Austria, June 1995. AACE.

[21] Ed H. Chi. Webspace Visualization. Available at `http://www.geom.umn.edu/docs/weboogl/`, The Geometry Center, University of Minnesota.

[22] Sougata Mukherjea and James D. Foley. Visualizing the World-Wide Web with the Navigational View Builder. *Computer Networks and ISDN Systems*, 27(6):1075–1087, April 1995. Proc. Third International World-Wide Web Conference, WWW'95, Darmstadt, Germany.

[23] Andrew Wood, Nick Drew, Russell Beale, and Bob Hendley. HyperSpace: Web Browsing with Visualisation. Available at `http://www.cs.bham.ac.uk/amw/hyperspace/`. Poster/Demo at WWW'95, Darmstadt, Germany.