Constructing Cyberspace: Virtual Reality and Hypermedia

Keith Andrews

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

Abstract

Large-scale, distributed hypermedia information systems allow fast, structured access to very large, dynamic information bases. The highly perceptual nature of a virtual reality interface has the power to take users both inside information and inside its structure. Combining the two takes us a step towards cyberspace, William Gibson's vision of a virtual model of all the world's interconnected data. This paper reviews current work on the boundary of virtual reality and hypermedia.

1 Introduction

Cross-fertilisation between virtual reality (VR) and hypermedia has resulted in two main application domains: "hypertextualising the space" and "spatialising the hypertext", as Jay Bolter termed them in his ECHT'92 keynote address [Bol92]. The former entails the use of navigable, three-dimensional models, certain elements of which embody hypermedia links to other related information. The latter involves generating a spatial representation of the hypermedia web itself, a kind of 3D overview map.

The use of navigable, three-dimensional models in hypermedia, though a seemingly natural extension of two-dimensional raster images or drawings, has yet to receive much attention in the hypermedia community. Jay Bolter envisages an immersive approach, "writing on the world", where textual and other information is portrayed directly in the virtual model in the form of posters, stone inscriptions, and the like. Links might be represented as some kind of flying device like a magic carpet, which whisks the reader to another location. Or maybe as a tunnel entrance, the tunnel leading to a new location, or a "magic window" into another world (which has the advantage of giving some kind of preview as to where one would go).

Here in Graz, we have implemented a "through-the-window" approach, where three-dimensional models are considered to be hypermedia documents in their own right and are viewed through a window on the computer screen; pieces of textual or other information being presented in windows of their own.

Graphical browsers and overview maps for hypermedia systems already exist in two dimensions, although their automatic generation (as opposed to hand-crafted) poses non-trivial problems, particularly in the area of automatic graph layout. Work on extending the idea to three dimensions is still at the research stage. However, the possibilities offered are very attractive: for example a 3D overview map where pieces of information are represented as shapes in 3-space, with links strung as cables between them and clusters representing related data.

2 Hypermedia and Deep Hyperspace

Hypertext is non-linear text. Unlike the typical printed book, which is read sequentially from beginning to end, a hypertext comprises many interlinked chunks of text. Readers are not bound to a particular sequence, but can browse through the information naturally by association, following their interests by clicking on highlighted keywords in one text to bring up another associated chunk of text. Hypermedia is the generalisation of hypertext to include all kinds of media (images, drawings, sound, film clips, 3D models, etc.) as well as text. The individual chunks of information are usually referred to as *documents* or *nodes*, and the connections between them as *links*. [Nie90] and [BD91] are good introductory texts on hypertext and hypermedia.

One of the main problems in hypermedia systems is that of user disorientation, becoming "lost in hyperspace", first discussed by [EH89]. The symptoms are: the difficulty of gaining an overview, finding information again, not knowing how much information there is on a subject and how much of it has already been seen, not knowing whether everything important has been seen, etc. To help users orient themselves in the information space, most hypermedia systems provide one or more navigational aids:

- **Guided Tours**. Predefined (mostly linear) paths through hyperspace.
- **Search Facilities**. Keyword search, full text search, fuzzy search (with errors), similarity measures [Sal89].
- **Visual Overviews**. Manually crafted drawings, maps, annotated photographs, etc.
- **Graphical Browsers**. Automatically generated structure maps [UY89]:
  - global map (entire hyperspace)
  - local map (around current node)
  - fisheye view (magnifying glass around current focus)
  - automatic clustering
  - hierarchical browser.
- **Backtracking**. Return to previous node, history list, bookmarks.

Most current hypermedia systems are frame-based, stand-alone, single-user systems (e.g. HyperCard [Goo87] and ToolBook [Cor89]). Documents must fit into a fixed-size frame (screen) and authors have full control over both their content and presentation.

Hypermedia information systems, on the other hand, are typically window-based, distributed, multi-user systems (e.g. Intermedia [HKRC92], Hyper-G [KMS93, KM93], and WorldWideWeb [BCGP92]). Documents may be any size, each displayed in its own (scrollable) window. Authors specify only the information content of documents, style sheets control their presentation (see [AK93, SKA93]).

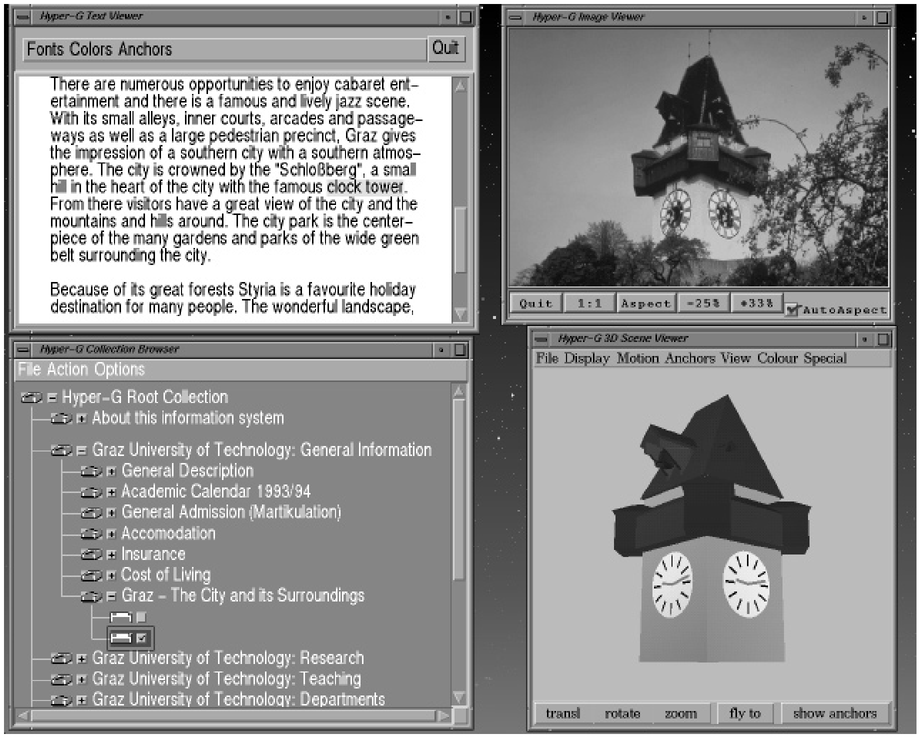

Figure 1: Harmony: The Hyper-G Browser for X Windows

Hyper-G is an ambitious, large-scale hypermedia project currently underway at Graz University of Technology. Figure 1 shows Harmony [AKJ94], the Hyper-G browser for X Windows. Harmony's collection (hierarchical) browser, text viewer, image viewer, and 3D viewer are visible.

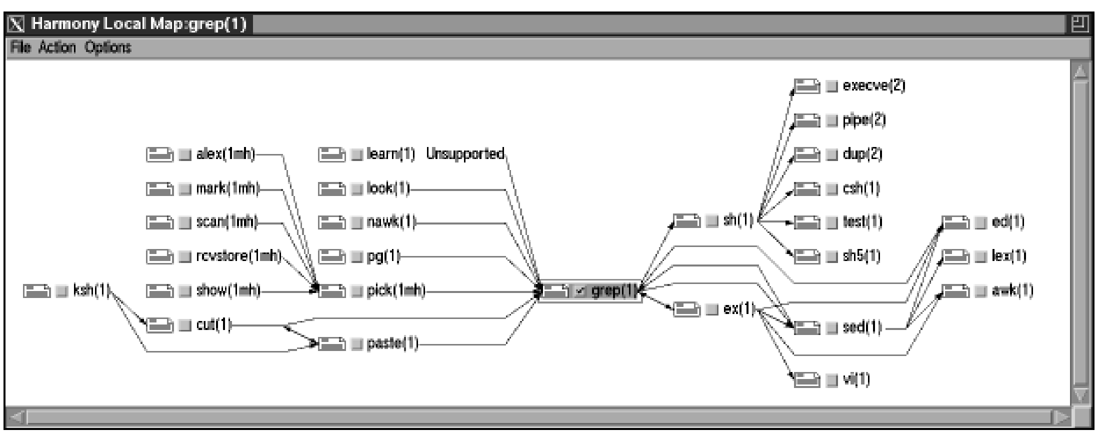

Figure 2: Harmony: Local Map Feature

Figure 2 shows a Harmony local map for the “grep” Unix manual page.

We use the term *deep hyperspace* to refer to large-scale (many tens of thousands of nodes), dynamic hyperstructures. The information base is in a permanent state of flux as documents and links are constantly updated, inserted, and deleted. Information is typically distributed over many sites. Deep hyperspace approaches William Gibson's *cyberspace*, a virtual model of all the world's interconnected data.

3 3D Overviews – Mapping Out Cyberspace

Some seminal work has been done on "spatialising the hypertext", providing three-dimensional overview maps and visualisations of (hypertextual) information.

SemNet [FPF88] 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 [Fur86] were also implemented.

The Information Visualizer [RCM93] provides three-dimensional representations for linear and hierarchically structured information: the perspective wall [MRC91] and cone tree [RMC91] 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 [FSHH93] 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).

Smith and Wilson [SW93] describe a test system based on HyperCard and Virtus Walkthrough (a 3D visualisation system). In the context of an academic departmental information system, they investigated both two and three dimensional schematic and spatial representations of the hypermedia network. The 3D schematic representation is similar to a cone tree. The 3D spatial representation involves a model of the department, with HyperCard cards dynamically presented as the user's location in the model changes.

The File System Navigator (FSN, or "Fusion") written by Joel Tesler and Steve Strasnick at Silicon Graphics [TS92] 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.

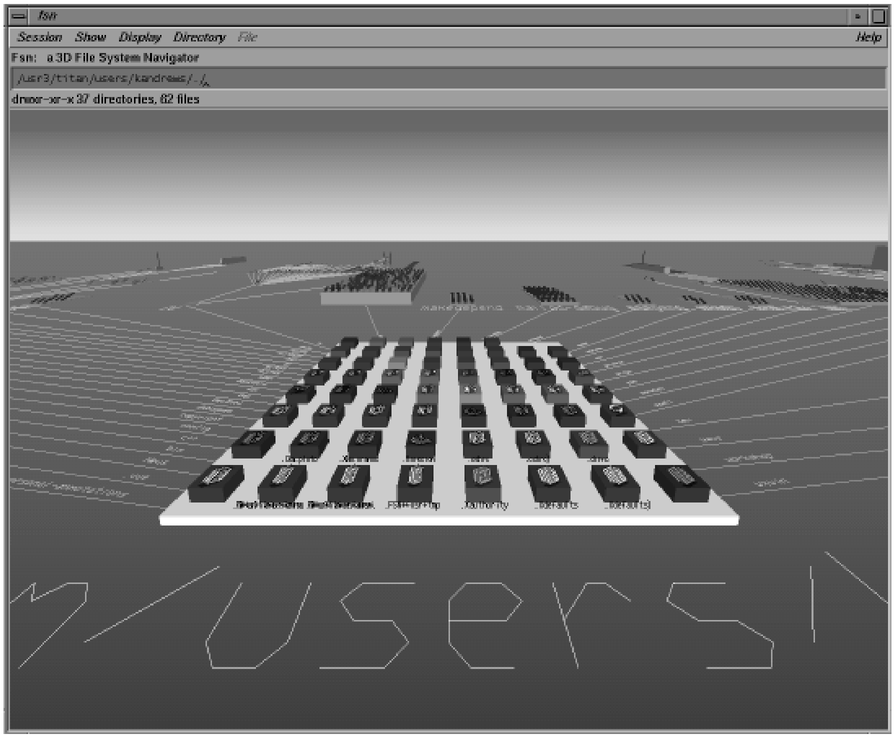

Figure 3: Fusion: The File System Navigator

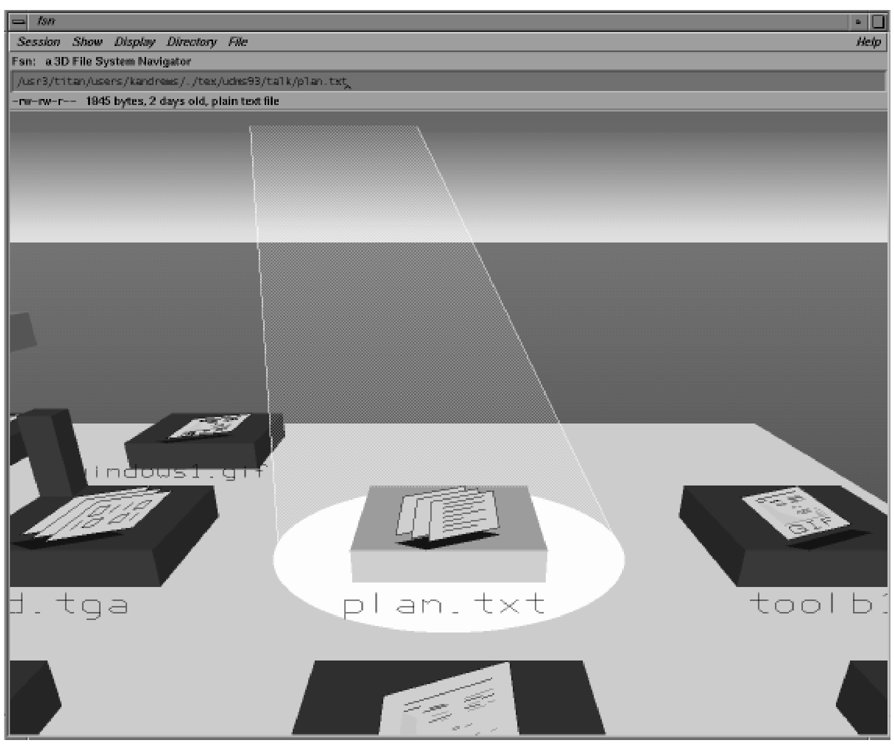

Figure 4: Fusion's Virtual Spotlight

Figure 3 shows a Fusion visualisation of my home directory and Figure 4 a text file highlighted by the virtual spotlight.

We are just starting a project in Graz to look at providing 3D overview maps for our Hyper-G hypermedia system.

4 3D Hyperdocuments

As for "hypertextualising the space", using three-dimensional models in a hypertextual context, things are just getting started. In his ECHT'92 keynote address, Jay Bolter presented some early work on his immersive approach, "writing on the world", where textual and other information is portrayed directly in the virtual model in the form of posters, stone inscriptions, and the like.

The Information Visualiser uses a 3D rooms metaphor. Particular data sets (nodes) are visualised in rooms of their own. Doors (links) set in the walls of a room can be opened (clicked) to reach another room through a kind of teleportation.

In Graz, we have implemented a fully-integrated 3D scene viewer as part of the Harmony browser for the Hyper-G hypermedia information system [And92]. Scene description files are stored in the hypermedia database just like any other document. Users arriving at such a scene node (by following a link from some other node in the hypermedia web) are presented with a 3D representation of the scene. They are then free to look around the scene interactively at their own leisure using a variety of navigational metaphors. Hypermedia links are anchored to individual objects within the scene; these may then be clicked to activate the link. Such anchor objects may be selectively highlighted, to gain an impression of which objects have links attached.

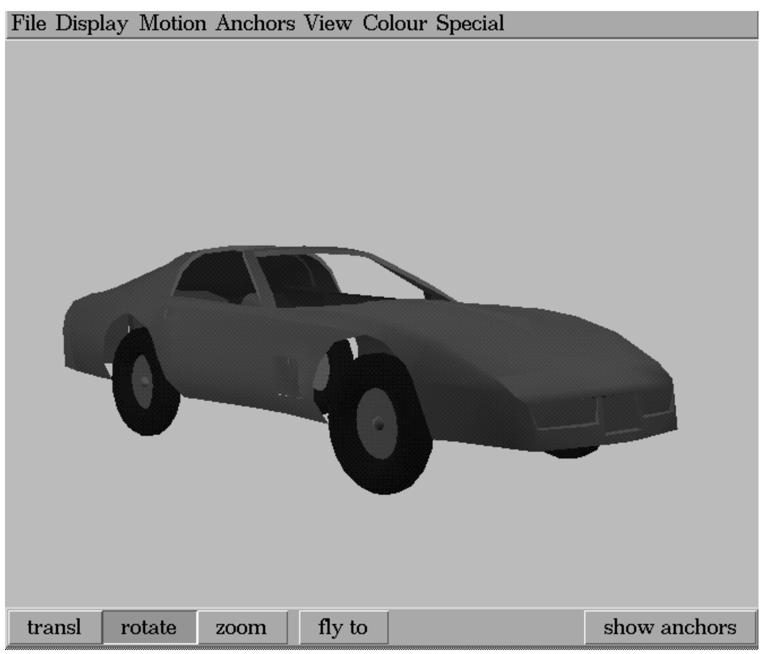

Figure 5: Harmony: 3D Viewer Displaying Model of Motor Car

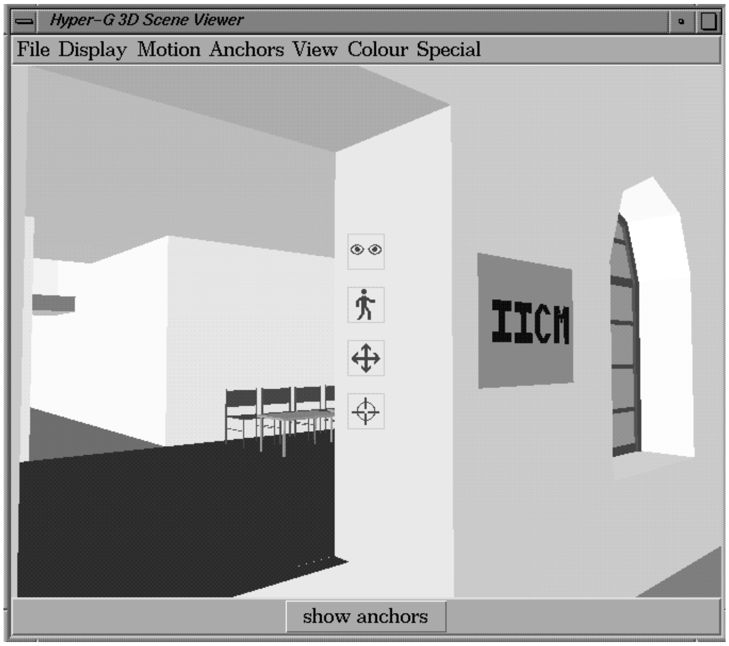

Figure 6: Harmony: 3D Viewer Displaying Model of IICM Buidling

Figure 5 shows the 3D viewer displaying a model of a motor car, as might be used in a hypermedia car handbook with 3D illustrations. The model may be translated, rotated, and zoomed, or the user can move in on a particular part of interest with the “fly to” metaphor (a variant of point-of-interest movement [MCR90]). Clicking on a tyre might open a text document about suitable types of tyre, etc.

Figure 6 shows navigation through a model of our institute building using a “heads-up” navigational metaphor similar to the one available in the Information Visualizer. Icons overlaid across the centre of the field of view provide quick access to specific controls. For example, pressing a mouse button within the “eyes” icon and dragging results in a turn of the head. Pressing within the “walk” icon and dragging results in motion forwards or backwards. The model contains numerous links to text, image, and scene documents.

5 Conclusion

I have given an overview of current work on the boundary of virtual reality and hypermedia. Two application domains are apparent: the provision of 3D overview maps, and the integration of 3D hyperdocuments. Regardless of whether an immersive or through-the-window approach is taken, elements of a VR interface have the potential to truly revolutionise access to hypermedia information systems in general, and to deep hyperspace in particular.

VR has until recently been perceived more as a playground for computer freaks than as a workplace for getting real work done. Combining elements of a virtual reality interface to a hypermedia information system is perhaps one way of bridging the divide and a step on the path towards the construction of cyberspace.

References

[AK93] Keith Andrews and Frank Kappe. Strait-jacketing authors: user interface consistency in large-scale hypermedia systems. In *Proc. Hypermedia '93, Zurich, Switzerland*, Springer-Verlag, Berlin, March 1993.

[AKJ94] Keith Andrews, Frank Kappe, and Jürgen Schipflinger. Harmony: a tool for navigating through deep hyperspace. 1994. Submitted to ED-MEDIA 94, Vancouver, Canada.

[And92] Keith Andrews. Using 3d scenes as hypermedia nodes. In *Poster at ECHT'92, Milan, Italy*, December 1992.

[BCGP92] Tim Berners-Lee, Robert Cailliau, Jean-François Groff, and Bernd Pollermann. WorldWideWeb: the information universe. *Electronic Networking: Research, Applications and Policy*, 2(1):52–58, Spring 1992.

[BD91] Emily Berk and Joseph Devlin, editors. *Hypertext/Hypermedia Handbook. Software Engineering Series*, McGraw-Hill, New York, 1991.

[Bol92] Jay David Bolter. Virtual reality and the future of hypertext. In *Proc. ECHT'92*, page 2, ACM Press, December 1992.

[Cor89] Asymetrix Corporation. *Using ToolBook. A Guide to Building and Working with Books.* Asymetrix Corporation, Bellevue, WA, 1989.

[EH89] Deborah M. Edwards and Lynda Hardman. Lost in hyperspace: cognitive mapping and navigation in a hypertext environment. In R. McAleese, editor, *Hypertext: Theory into Practice*, pages 105–125, Blackwell Scientific Publications Ltd., 1989.

[FPF88] 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.

[FSHH93] Kim Michael Fairchild, Luis Serra, Ng Hern, Lee Beng Hai, and Ang Tin Leong. Dynamic fisheye information visualizations. In Alan Wexelblat, editor, *Virtual Reality: Applications and Explorations*, pages 45–74, Academic Press, 1993.

[Fur86] George W. Furnas. Generalized fisheye views. In *Proc. CHI'86, Boston*, pages 16–23, ACM, April 1986.

[Goo87] Danny Goodman. *The Complete HyperCard Handbook.* Bantam, New York, 1987.

[HKRC92] Bernard J. Haan, Paul Kahn, Victor A. Riley, James H. Coombs, and Norman K. Meyrowitz. IRIS hypermedia services. *Communications of the ACM*, 35(1):36–51, January 1992.

[KM93] Frank Kappe and Hermann Maurer. Hyper-G: a large universal hypermedia system and some spin-offs. *ACM Computer Graphics, experimental special online issue*, May 1993. Available by anonymous ftp from siggraph.org in directory publications/May_93_online/Kappe.Maurer.

[KMS93] Frank Kappe, Hermann Maurer, and Nick Sherbakov. Hyper-G – a universal hypermedia system. *Journal of Educational Multimedia and Hypermedia*, 2(1):39–

66, 1993.

[MCR90] Jock D. Mackinlay, Stuart K. Card, and George G. Robertson. Rapid controlled movement through a virtual 3D workspace. In *Proc. SIGGRAPH '90, Dallas, Texas*, pages 171–176, ACM, New York, August 1990.

[MRC91] 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, ACM, New York, May 1991.

[Nie90] Jakob Nielsen. *Hypertext & Hypermedia*. Academic Press, San Diego, CA, 1990.

[RCM93] 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.

[RMC91] 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, ACM, New York, May 1991.

[Sal89] Gerard Salton. *Automatic Text Processing*. *Series in Computer Science*, Addison-Wesley, Reading, MA, 1989.

[SKA93] Robert Stubenrauch, Frank Kappe, and Keith Andrews. Large hypermedia systems: the end of the authoring era. In *Proc. ED-MEDIA 93, Orlando, Florida*, pages 495–502, AACE, Charlottesville, VA, June 1993.

[SW93] Pauline A. Smith and John R. Wilson. Navigating in hypertext through virtual environments. *Applied Ergonomics*, 24(4):271–278, August 1993. Butterworth-Heinemann Ltd.

[TS92] Joel Tesler and Steve Strasnick. *FSN: The 3D File System Navigator*. Silicon Graphics, Inc., Mountain View, CA, 1992. Available by anonymous ftp from sgi.com in directory fsn.

[UY89] Kenneth Utting and Nicole Yankelovich. Context and orientation in hypermedia networks. *ACM Transactions on Information Systems*, 7(1):58–84, January 1989.