To appear in Proc. of Eurographics Symposium and Workshop on Multimedia:
Multimedia/Hypermedia in Open Distributed Environments, Graz, Austria, June 1994.

Soaring through hyperspace: A snapshot of Hyper-G and its Harmony client

Keith Andrews and Frank Kappe

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

Abstract. This paper describes the current status of work on Hyper-G and its new viewer, Harmony. Hyper-G is a general-purpose, large-scale, distributed hypermedia information system under development at Graz University of Technology. It is based on the client-server model across the Internet and is interoperable with both Gopher and World Wide Web.

Harmony is the new native Hyper-G client for X Windows on Unix platforms. It takes advantage of Hyper-G's structuring and retrieval features to provide both intuitive navigational facilities and informative feedback about the location of information.

1 Introduction

The requirements of a large volume of information impose certain design decisions on the implementation of a large-scale hypermedia information system. An important issue is support for automatic structuring and maintenance of a dynamically changing body of information. Another aspect of the size of hypermedia datasets is that orientation and navigation become more difficult as size increases. Problems of users of such systems include: becoming "lost in hyperspace", having difficulty gaining an overview, not being able to find information that is known to exist, determining how much information on a given topic exists, how much of it has been seen, and how much is left. These issues have been identified as crucial for the acceptance of hypermedia and have been intensely discussed in the literature (see for example [5, 13, 14]). However, solutions which work well on small systems fail completely when applied to large-scale hypermedia (such as global maps [20]). The Hyper-G project was started in 1990 in order to explore these issues as applied to a large-scale, distributed hypermedia information system.

The first Hyper-G client was a simple terminal viewer[1], which is now a stable, universal point of entry to any Hyper-G server. In October 1992 work started

[1] The terminal viewer can be tried out by `telnet`'ing to `hyperg.tu-graz.ac.at` which starts a session on Graz University of Technology's main Hyper-G server.

on Harmony, a Hyper-G client for X Windows on Unix platforms. Harmony is a modern, Motif-style graphical interface for Hyper-G with Version 1.0 scheduled for release in mid-1994.

This paper presents a snapshot of the current status of Hyper-G and Harmony, emphasising particularly interesting and original features not found in comparable systems.

2 Hyper-G

Hyper-G is designed as a general-purpose, large-scale, distributed, multi-user, hypermedia information system, similar in scope to Xanadu [12], Intermedia [8], WAIS [18], Gopher [1], and World Wide Web (WWW) [3]. Based on previous experience with large-scale videotex information systems, the aim of Hyper-G is to develop a flexible hypermedia framework in order to study and possibly eliminate the problems typically associated with large-scale hypermedia systems. The basic concepts underlying Hyper-G have been presented elsewhere [11, 10, 9], here we provide a general overview of current functionality.

At the server level, Hyper-G provides a number of orthogonal structuring and retrieval facilities:

- Structuring of documents into so-called *collections*, which may themselves belong to other collections. Navigation may be performed down through the collection hierarchy, access rights assigned on a collection-by-collection basis, and searches restricted to particular collections.
- Hyperlinks from a *source anchor* within one document to either a *destination anchor* within another document, an entire document, or a collection. Links are not stored within documents (as in WWW) but in a separate database, which has the advantage that they are bidirectional and may be added to read-only documents.
- Attribute and full text search. Documents have an associated set of attributes (author, title, keywords, etc.) which may be searched for, including boolean combinations. Full text search facilities include fuzzy boolean queries [16] and WAIS-like nearest-neighbour searches based on the vector space model [17]. Every document and collection is automatically added to the full text index on insertion into the database. The scope of a search may be as narrow as one collection on a single server or as wide as all collections on all Hyper-G servers worldwide.

Other features supported by Hyper-G and not found in comparable systems such as Gopher and WWW include:

- Four user identification modes: from anonymous to fully identified.
- Support for user groups.
- Access rights per user group for documents and collections.

- Support for multilingual versions of documents and language preferences.
- An underlying object-oriented database to guarantee the consistency and integrity of data (for example the updating of links when a document is moved or the elimination of dangling links when a document is deleted).

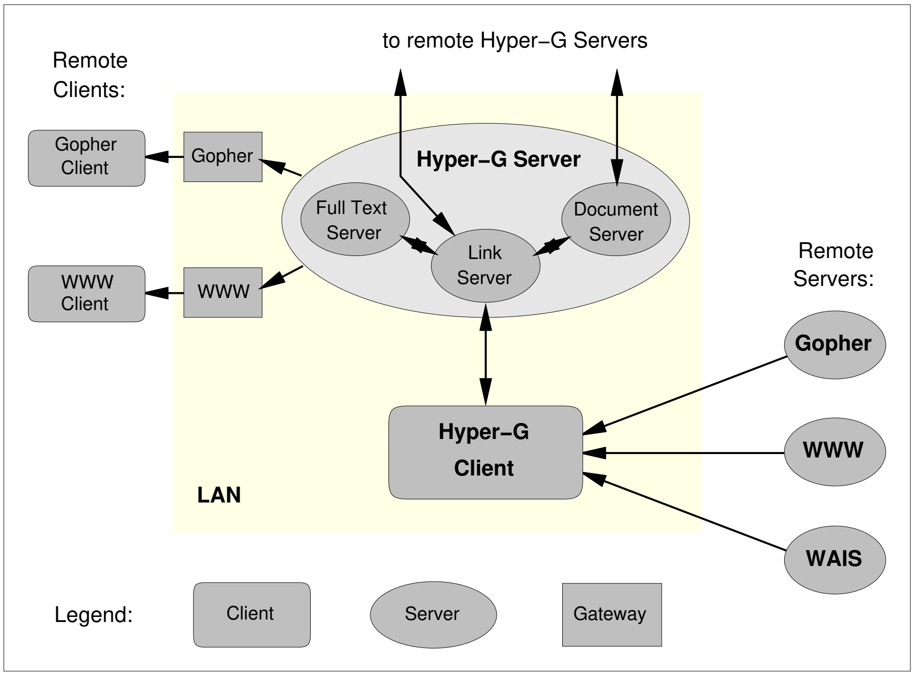

Fig. 1.: The architecture of Hyper-G

Hyper-G is based on a client-server model across the Internet. Figure 1 shows the architecture of Hyper-G. Unlike Gopher or WWW clients which connect to many servers during a typical session, Hyper-G clients talk to a single Hyper-G server for the entire session. Should information from a remote server be needed, the local server fetches it and passes it on to the client. This approach has the following advantages:

- Clients are kept simple.
- An efficient, connection-oriented protocol can be used.
- Remote information can be cached in the local server.
- User accounts and access rights have only to be maintained in the local server (the user has to identify to one server only).
- Statistics and user profile information can be gathered on a per-session basis.

Also apparent from Figure 1 is the interoperability of Hyper-G with Gopher and WWW clients and Gopher, WWW, and WAIS servers. When accessed by a Gopher client, the Hyper-G server maps the collection hierarchy into a Gopher menu tree (hyperlinks cannot be represented in Gopher). A synthetic search item is generated at the foot of each Gopher menu to allow searching the corresponding collection. When accessed by a WWW client, each level of the collection hierarchy is converted to an HTML [4] document containing a menu of links to other sub-menus. The menus are marked as searchable. Hyper-G text documents are transformed on-the-fly into HTML documents, including any links they might have.

In the other direction (the right hand side of Figure 1), Hyper-G clients can contact Gopher, WWW, and WAIS servers in order to retrieve information from them. The Hyper-G server is able to store pointers to such remote objects. This allows the incorporation of information on remote non-Hyper-G servers (almost) seamlessly: Gopher menus are transformed into Hyper-G collections, WWW text documents into Hyper-G text documents, and WAIS queries and responses into Hyper-G queries and responses. We are in the process of moving the knowledge of external protocols from Hyper-G clients to the Hyper-G server, in order to make clients simpler and enable caching of external documents.

3 Harmony

Harmony is the native Hyper-G client for X Windows on Unix platforms. It takes advantage of Hyper-G's structuring and retrieval features to provide both intuitive navigational facilities and informative feedback about the location of information. A number of interesting and novel features, including several kinds of dynamic overview and hierarchy maps, a three-dimensional scene viewer, and three-dimensional navigational aids, have been implemented.

Harmony is a multi-process Unix application (see Figure 2), written in C++, using the InterViews X11 user interface toolkit, and (for its 3D features) Silicon Graphics' GL graphics library. The primary process is the *session manager*, which communicates with the Hyper-G server, provides navigational facilities, and coordinates all other activities. The session manager starts secondary processes, so-called *document viewers*, as necessary to display particular documents.

Native Harmony viewers for text, images, 3D scenes, and MPEG films are currently available. External applications are started to display other document types, such as audio clips and PostScript files. Harmony may optionally be configured to run external programs instead of any native viewer (the document is piped to standard input), but with the restriction that link activation and editing are no longer possible.

Figure 3 shows an example Harmony session. The session manager (left), text viewer (right) and image viewer are visible. Note how the session manager provides for navigation through the collection hierarchy. Collections may be

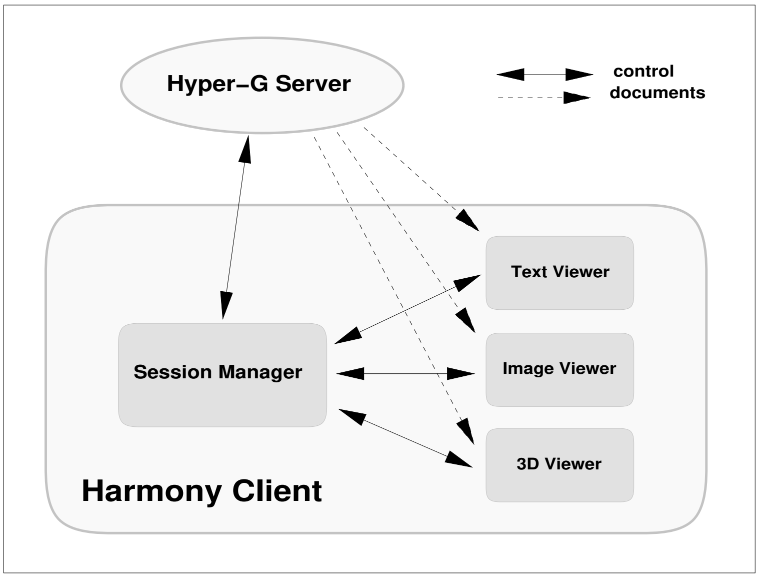

Fig. 2.: The architecture of Harmony

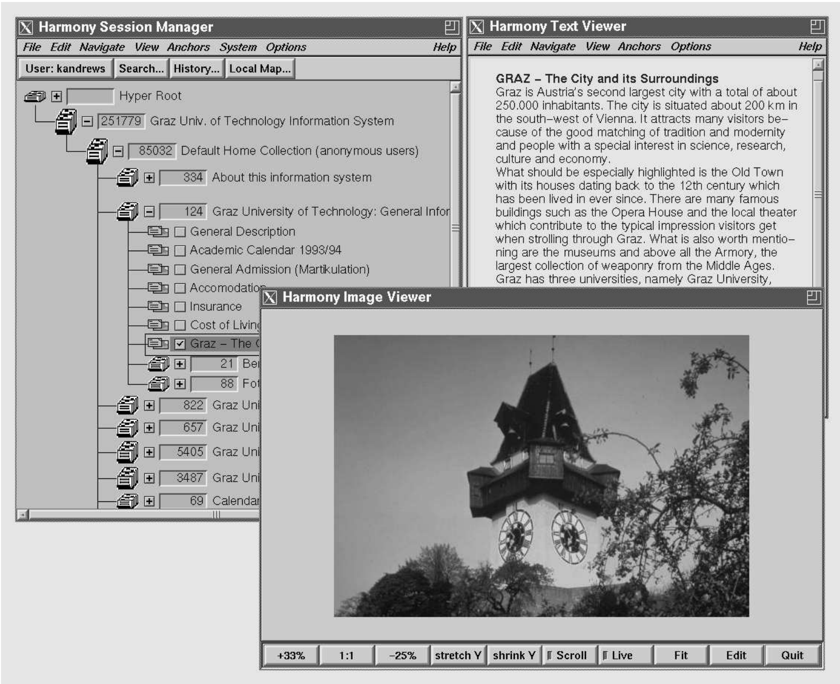

Fig. 3.: Example Harmony session

opened and closed and documents activated by double-clicking. Collections or documents which have already been visited are marked with a tick. In this example, a descriptive text and an image about the city of Graz have been accessed.

The Harmony text viewer displays marked-up (SGML) text documents, and has the usual facilities for scrolling, searching, selecting, etc. The display styles of the various attributes (title, bold, anchor, etc.) are user-configurable. Link activation is accomplished by double-clicking a text anchor and implemented by passing a message back to the session manager.

The image viewer accepts raster images in a variety of common formats (TIFF, GIF, etc.). Common operations such as zooming and panning are available. Link anchors in an image may be of a variety of formats: polygonal areas upon the image, buttons positioned at a particular location, etc.

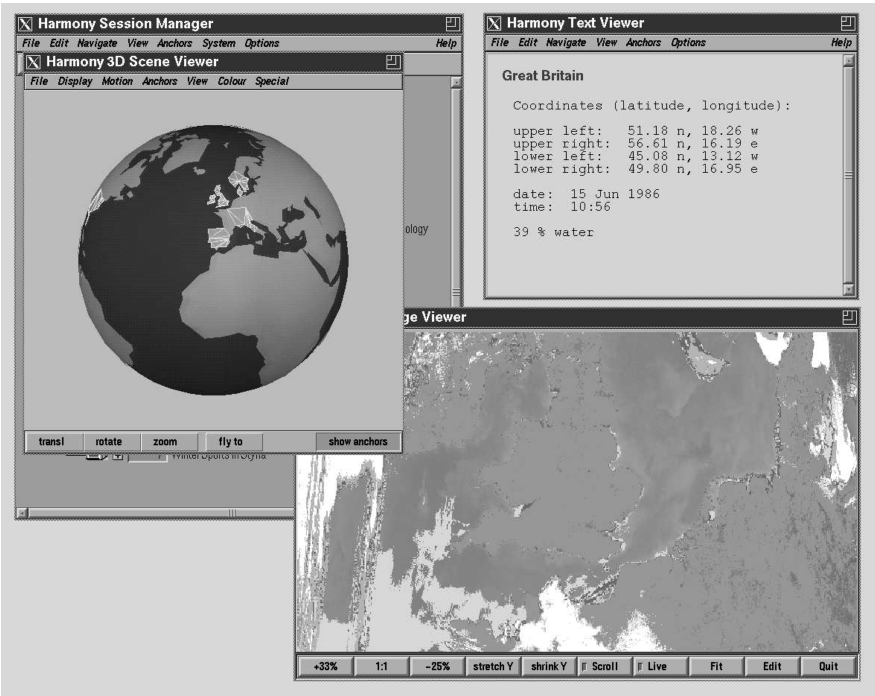

Fig. 4.: Harmony scene viewer

A unique feature of Harmony is its support for 3D scene documents. Scene description files representing arbitrarily complex three-dimensional models of scenes or objects are displayed by the Harmony 3D scene viewer. Figure 4 shows the scene viewer displaying a model of the globe. The anchor representing Great Britain has been clicked to retrieve the corresponding satellite image. Users can manipulate an object (translate, rotate, zoom) and navigate (walk, fly, fly to, heads-up) within a scene, in a fashion similar to the Information Visualiser [15]. Hypermedia links may be attached either to individual objects within a scene or to groups of polygons within an object.

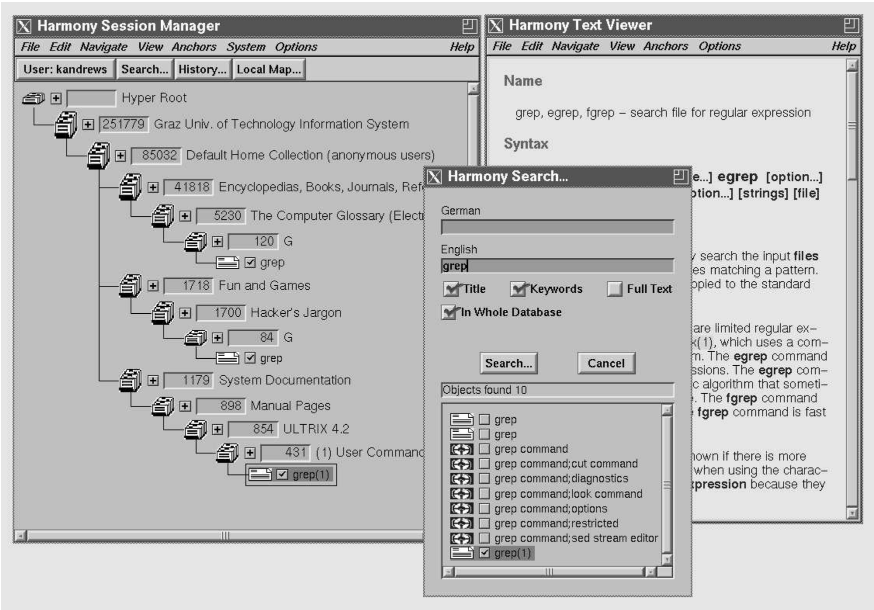

Fig. 5.: Harmony search dialogue

Harmony supports title, keyword, and full text search. The search is performed in the current collection by default. The output is an ordered list of matching documents. This list is of course active, so that documents may be activated by double-clicking within the search output list. Figure 5 shows a search dialogue for the word "grep" in the titles and keywords of all documents and anchors on the local server. Three text documents and a number of anchor objects have been found. Clicking on the three text documents in turn has caused the session manager to open up the path to each of them in the collection hierarchy. Having seen that the bottom-most document belongs to the collection "User Commands" in collection "ULTRIX 4.2" in "Manual Pages", the user has activated this document in preference to the Hacker's Jargon or Computer Glossary entries. This *location feedback* is an important feature of Harmony – it enables users to make intelligent choices about search results before committing to fetch a particular document. Furthermore, it allows users to build up a mental picture of the locations of documents.

The *local map* facility, similar to the local map of Intermedia, provides a kind of short-range radar, generating on request (dynamically) a map of the vicinity of a document. By default, two levels of incoming and outgoing hyperlinks are represented. Figure 6 shows the local map for the "grep" manual page. Of

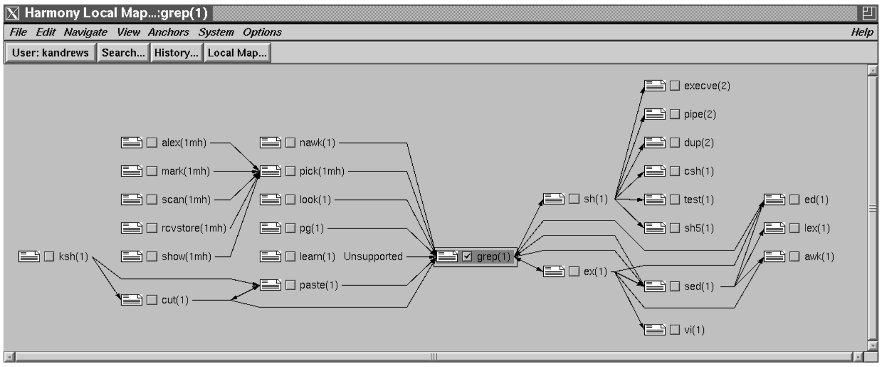

Fig. 6.: Harmony local map

course the local map is active; documents may be activated by double-clicking their icons.

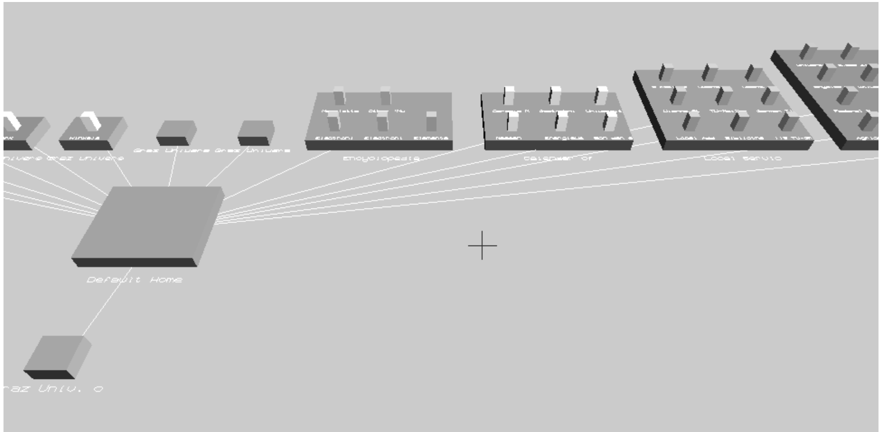

Fig. 7.: Harmony information landscape

Work is currently under way on a three-dimensional graphical overview map, called the *information landscape*, a prototype of which is shown in Figure 7. The collection hierarchy is mapped out onto a plane and the third dimension used to encode size. Users can "fly" over the hyperspace landscape looking for salient features, like flying over a file system with FSN [19]. The information landscape is an alternative to navigating with the session manager's built-in collection hierarchy browser. In the coming months, truly three-dimensional layouts will be investigated [6, 7] and work towards an immersive (virtual reality) interface is planned [2].

Authoring features currently under implementation in Harmony include a drag-and-drop interface to the local file system, allowing documents to be simply pulled into collections in the session manager, and a link editing facility supported across all Harmony viewers.

4 Concluding Remarks

We have pesented a snapshot of the current status of Hyper-G and its Harmony client for X Windows. More detailed information and a number of technical reports are available by anonymous ftp from `iicm.tu-graz.ac.at` in directory `pub/Hyper-G/papers`.

5 Acknowledgements

Partial Support of the Hyper-G project by the Austrian Ministry of Science, Joanneum Research, and the European Space Agency is gratefully acknowledged.

References

[1] Bob Alberti, Farhad Anklesaria, Paul Lindner, Mark McCahill, and Daniel Torrey. The internet gopher protocol: a distributed document search and retrieval protocol. March 1992. Available by anonymous ftp from boombox.micro.umn.edu in directory pub/gopher/gopher_protocol.

[2] Keith Andrews. Constructing cyberspace: virtual reality and hypermedia. December 1993. To appear in Proc. of Virtual Reality Vienna '93.

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

[4] Tim Berners-Lee and Dan Conolly. Hypertext markup language (HTML). June 1993. Version 1.2. Available in hypertext on the World Wide Web as http://info.cern.ch/hypertext/WWW/Markup/HTML.html.

[5] 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.

[6] Kim Michael Fairchild. Information management using virtual reality-based visualizations. In Alan Wexelblat, editor, *Virtual Reality: Applications and Explorations*, pages 45–74, Academic Press, 1993.

[7] 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.

[8] 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.

[9] Frank Kappe. Hyper-G: a distributed hypermedia system. In Barry Leiner, editor, *Proc. INET '93, San Francisco, California*, pages DCC–1–DCC–9, Internet Society, August 1993.

[10] 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.

[11] Frank Kappe, Hermann Maurer, and Nick Sherbakov. Hyper-G – a universal hypermedia system. *Journal of Educational Multimedia and Hypermedia*, 2(1):39–66, 1993.

[12] Theodor Holm Nelson. *Literary Machines (Edition 87.1)*. The Distributors, South Bend, IN 46618, USA, 1987.

[13] Jakob Nielsen. The art of navigating through hypertext. *Communications of the ACM*, 33(3):296–310, March 1990.

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

[15] 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.

[16] Gerard Salton, E. A. Fox, and H. Wu. Extended boolean information retrieval. *Communications of the ACM*, 26(12):1022–1036, December 1983.

[17] Gerard Salton, A. Wong, and C. S. Yang. A vector space model for automatic indexing. *Communications of the ACM*, 18(11):613 ff., November 1975.

[18] Richard Marlon Stein. Browsing through terabytes – wide-area information servers open a new frontier in personal and corporate information services. *Byte*, 16(5):157–164, May 1991.

[19] 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.

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